


CONNECTIONS



The Interoperability Report

October 1993

Volume 7, No. 10

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Live from INTEROP.....	2
Connecting to the Internet.....	6
Establishing Workgroups.....	10
Associations.....	19
Letters to the Editor.....	22
Book Review.....	24
Announcements.....	25

From the Editor

INTEROP 93 August is over. I have just enough time to publish an “in-between” edition before preparing the last INTEROP companion issue for the year—the Paris conference and exhibition takes place October 25–29 and our next issue (November) will focus on networking from a European perspective.

But back to this month: Our first article is a report from INTEROP by Jack Kessler. He actually sent the text to me during the show, using a wireless terminal. Wireless e-mail was only one of many new technologies on display at the show. With more than 65,000 attendees and 500+ exhibitors occupying every square foot of the convention center, INTEROP 93 August was clearly a great success.

The question I get asked most frequently these days is “how do I get on the Internet?” The answer depends, of course, on what you mean by “on the Internet,” so I usually go through the exercise of explaining the range of options from e-mail access to a full blown IP connection. Luckily, there are now books in the works (and some already in print) that explain all this, and we are happy to present an excerpt from the forthcoming book *The Internet Connection: A Guide to Connectivity and Configuration*, by John Quarterman and Smoot Carl-Mitchell.

Designers of corporate data networks are faced with a number of technical as well as administrative challenges. When isolated local area and wide area networks are brought together, one typically faces the “multiprotocol problem.” Bob Meindl outlines the establishment of *workgroups* as a solution, and explains the steps taken with the help of an example network.

Users, vendors and developers of high-speed networking technologies have formed *Associations* or *Interest Groups*, to foster standardization, development and deployment. Three such groups are *The ATM Forum*, *The Frame Relay Forum*, and the *SMDS Interest Group*. This month we bring you a brief overview of each organization and its activities, starting on page 19.

It’s been quite a while since we’ve had any “Letters to the Editor,” but this month we have two. Your comments, questions and suggestions are always welcome. Our e-mail address remains the same: connexions@interop.com

You’ll probably notice that the type in this edition looks crisper than in the past. This is because we just acquired a 600 dpi laser printer which will be used from now on for “camera ready copy.” It is interesting to note that we paid about one-third as much for this printer as we paid for our original 300 dpi printer which was purchased back in 1986. Technology marches on!

ConneXions is published monthly by Interop Company, 480 San Antonio Road, Mountain View, CA 94040-1219, USA. +1 415-941-3399. Fax: +1 415-949-1779. Toll-free: 1-800-INTEROP. E-mail: connexions@interop.com.

Copyright © 1993 by Interop Company. Quotation with attribution encouraged. *ConneXions—The Interoperability Report* and the *ConneXions* logo are registered trademarks of Interop Company.

ISSN 0894-5926

Live from INTEROP:
TV/phone/computer/Internet mergers

by Jack Kessler, AKCO Inc.

Introduction

The following conference report is being transmitted on a wireless radio modem-equipped PowerBook laptop, via a little "cigarette-girl" front-pack display rack, carried by a guy named Chris who is part of a team walking around San Francisco's Moscone Convention Center, selling this service (from RadioMail Corporation) to INTEROP Conference attendees:

INTEROP ("interoperability") is where the computer world now goes to find out what is coming next. The conference is only a few years old but has grown wildly, so much so that today it is one of the largest conventions hosted by San Francisco: all the hotels here are full tonight, and rumors are that conference attendance will top 65,000.

Computer people come to INTEROP because "interoperability"—networks, and linking different kinds of electronic gizmos to each other—apparently is where they have decided their future is "at." Television (including cable!) and telephone people are here too, presumably for the same reason. Everyone looks a little worried—they all have well-publicized financial troubles—but they are excited as well: they seem to think that INTEROP offers the next big hi-tech solution. I've been worried as well, as a librarian and information-user: I still can't find things easily on the Internet, and every time the "techies" tweak the system it gets more complicated and much, much bigger—I wanted to find out what they think is in store. I have, I think.

**The Deaths of TV and
the Telephone**

The keynote speaker Wednesday morning declared the deaths of television and the telephone. George Gilder, "futurist," ex- of Harvard Economics and Henry Kissinger's staff, declared that digital cellular telephony coupled with wireless computer networks will replace TVs, telephones, and computers, all as we now know and love or hate them, within 20 years. Gilder also said some other very interesting things. "A computer without a network is like a car in the jungle," for example: the one without roads or service stations, the other with nothing to plug into to "get online"—interesting to consider how important networks have become, in such a short time, to the little boxes which we used to call our computing capacity. Gilder described the three technologies of "sand" (silicon chips), "glass" (fiber-optics), and "air" (the electromagnetic spectrum) which are emerging to rule the new "telecosm": he strongly contends that the technology of "air" is infinite—that there will be enough bandwidth to go around for cellular technologies. He also called the Clinton/Gore *National Information Infrastructure* (NII) effort a "cock-a-doodle-doo" policy, an effort to "celebrate the sunrise": the sun comes up anyway, Gilder said, and the "telecosm" is going to land on us whether Washington develops its NII or not. The "death of TV, the telephone and the computer" prediction, though, was by far Gilder's most dramatic pronouncement.

Mark Twain once protested that rumors of his death had been "greatly exaggerated." Twain's jibe has plagued "futurists" ever since. Still, predictions can be fun, and they can serve to animate and inspire an audience. Gilder's woke up our 9am audience. The part which woke me up most was the "20 years": setting a time-limit is the most dangerous part of "futurology"—but now I have encountered Chris and his front-pack wireless walking Internet workstation, and I am becoming somewhat of a believer.

Conference sessions for humans and Non-

INTEROP conference events number in the many dozens: 72 *Sessions* attended by 200–500 people each, 41 pre-conference *Tutorials*, various *Executive Sessions*, *Special Sessions* and *Plenary Sessions*, *Birds Of a Feather* (BOF) meetings which seem to go until the middle of the night, covering a vast variety of issues, “invitation only” vendor-sponsored parties at the hotels, amazing amounts of intense corridor-conversation and politicking and selling.

INTEROP sessions have titles like “SNMPv2,” “Implementing Distributed Object Services,” and “Alien Protocols (TCP/IP and OSI Networks) in SNA”; although there are the more human-sounding “Public Data Networks,” “Managing the Management Process,” and “How to Write the Best RFP,” as well. Clifford Lynch, Paul Peters, Jim Fullton and Howard Besser are here, presenting, “Online Libraries, Electronic Libraries and Networked Information Resources.” One has to choose.

Computer systems and users are converging, they say (both of them), which interests me, and I am not too good at things which call themselves names like “SNMPv2,” and I have a general allergy to acronyms, so I’ve decided to go to anything in human language having to do with users.

So far I’ve seen an excellent presentation by Susan Estrada—ex- of CERFnet, now of Aldea Communications—called “Internet 101,” then an interesting grope toward users by a group calling themselves “Commercial Use of the Internet,” and finally an even more interesting although somewhat distressing session ably managed by author Ed Krol on “Personal Access to the Internet.”

“Internet 101”—the basic course

Susan Estrada’s basic take on the Internet is that a “can-do” approach will work, for just about anything which one wants to do, if one is realistic. Her “can-do” approach came through in discussions, during her two-day tutorial marathon, of the Internet’s infamous *Acceptable Use Policies* (AUP). Susan basically considers such policies worthless, and frankly describes the ease with which countless efforts have succeeded in circumventing them since the Internet’s inception. Listening to Susan, one gets the impression that those of us who have been respecting AUP restrictions on non-academic network use have been in a very small minority for a very long time. Academic AUP does appear to be dying at last this year—the Internet is going public, and commercial—but her description of the historical holes in AUP can serve as an eloquent warning to future government policy-makers against narrow-minded, restrictive, and most of all short-sighted networking policies.

Susan’s realism came through in her basically negative assessment of accessibility on the Internet. She presented most of the new Internet “tools” in her tutorial, but was frank and often funny in her description of her own difficulties in getting these tools—*mail*, *FTP*, *Telnet*, *USENET*, *Gopher*, *WWW*, *WAIS*, *Archie*—to work right. It’s not a perfect world on the Internet yet, she told us, and it’s far better that we admit this and admit it to our users than that we pretend that we can do the impossible on what nevertheless still is the most exciting new information resource since the printed book.

Commercial use

The session on “Commercial Use of the Internet” featured three commercial users describing how a business can use this hitherto-academic resource efficiently, and two hopeful providers of the Internet resource to such commercial users. The users on the panel included an “instant” printer, someone from Silicon Graphics, and someone from a stock brokerage firm.

continued on next page

Live from INTEROP (*continued*)

Time value of information, ease of access of the technology, reliability of the network (the Internet is very, the other alternatives are not so), and security and confidentiality were much discussed. Great savings over the cost of leased lines appear to be a primary factor. At the close, John Curran of NEARnet warned that commercial users should, "be prepared to use the Internet or lose out to those who do."

Personal access

The session on "Personal Access to the Internet," chaired by author Krol, presented four individuals who currently are working on providing access to the general public: Martin Schoffstall of PSI, Edward Vielmetti of Msen, Walter Howe of DELPHI, and Michael O'Dell of UUNET. Krol cautioned that analogies mislead: the Internet is less like the "highway" which Al Gore likes to invoke, he said, than it is like airplanes taking off and landing from a vast system of airports—the real question now becoming, he added, what type of plane did the user take? Schoffstall described the efforts which his firm and others are making now to reach the general public with Internet services. He made the interesting suggestion that policy makers consider, "the environmental impact of removing one day of automobile commuting for all information workers in a metropolitan area." Vielmetti warned that public access still needs work. There still are many entry barriers, he said: primary among these is the lack of market research—the question still to be asked is, "what will motivate people at home to use this technology?" Howe proudly described his own firm's efforts: DELPHI is the 5th largest "information provider" now, he says—after CompuServe, Prodigy, Genie, and America Online—and in the last six months has added 40,000 new subscribers. O'Dell added his view that the "new transnational marketplace" for networking is showing a shift in emphasis, from Internet connectivity to Internet access.

I asked a question then. I hadn't heard much discussion yet, then or elsewhere at INTEROP I said, of the folks whom I consider to be "the general public." Walt Howe describes his users as including, "experienced school/work Internet users," "BBS users," "professionals looking for access," "new modem users," and, interestingly, "the visually impaired." I told the panel that these categories did not include "the general public," as I understand the term. I think more of someone very un-professional and in-experienced, who is, moreover, not "looking for" anything and more probably really doesn't want to be bothered. That's a fair description of the "general public" which many libraries would like to reach, and certainly of the "general public" which must be reached by any successful effort to "go public" with the Internet.

After some initial hemming and hawing I got some brave and, I think, honest answers. Walt Howe mentioned DELPHI's "talk"-style online interactive help, and their unique (?) *Gopher* capacity to combine FTP with zmodem to get data to a home computer: both are part of their conscious effort to make all this technology increasingly invisible, he said. Ed Krol was frank: we're not to the general public, thus defined, yet—but he is most encouraged by the success of many K-12 Internet pilot efforts, which he feels will create a "next generation" of network users. Ed Vielmetti thoughtfully mused, then and after the session, that a little "Internet diskette" in a pretty package, on the shelf at Egghead Software (for \$19.95)—one that would handle "login" and "password" and "billing" and "Z39.50" and "WWW" all "invisibly" for the user—might be a nice thing to have.

Internet Marketing— Calm before tidal wave?

I couldn't help thinking, though, as I toured through the vast commercial exhibits in the enormous Moscone Center later—over 500 vendors, names like IBM and Apple and AT&T and MCI and Pacific Bell and Klever Computers—that unless these folks get their marketing together quickly the “general public” may escape them. The really big sharks are gathering on this one, and they are really hungry, and some of them know marketing to the general public very well.

Ed Vielmetti and I had a talk over dinner and came up with a 3-part scenario (Ed agrees with most of this, I think): 1) the telephone companies will handle the “bits,” 2) the cable companies will handle the generic services (entertainment, mass media, the “sex 'n drugs 'n rock 'n roll”), and 3) the small rugged independents like the pioneers at INTEROP still will have the personalized services (training, navigating, filtering, anything lower-budget, higher-risk, labor-intensive, of which there still will be a lot under any realistic scenario). Library service, in my opinion, fits in somewhere in category #3. There is no guarantee, of course, that any of these three players will survive, particularly at the hands now of the other two: there are no rules written—the game still is being invented.

A comment of George Gilder's supports this scenario. “Intelligence will be on the periphery,” he asserts: the days of centralization are over, he says—the “pipe” will be neutral. This, if it's true, bodes well for both the large and the small players here: there will be a role for small, rugged independents, developing personalized networking, just as there will be a need for large economies-of-scale monoliths, which will supply and maintain the mass market, and the “pipe.”

One remaining question, though, is how to accommodate the very major change which appears to be implicit in the burst of activity going on this week at INTEROP. I'll be back at the conference during the next two days, poking into sessions, wandering the trade floor: my thought is that any librarian or information worker about to make a large hardware or network acquisition, or wondering about upgrading a telephone service, or puzzling over the future of a “media” center, or fussing with the continuous problems of logins, modems, syntaxes, norms and compatibility, might want to be at INTEROP too.

References

- [1] Pitsker, B. “Insights into the INTEROPnet,” *ConneXions*, Volume 7, No. 3, March 1993.
- [2] Curran, J. “The Internet as a tool for Business Communications?” *ConneXions*, Volume 7, No. 9, September 1993.
- [3] Deutsch, P. “The Internet as Service Provider,” *ConneXions*, Volume 7, No. 9, September 1993.
- [4] Cerf, V., “A National Information Infrastructure,” *ConneXions*, Volume 7, No. 6, June 1993.

JACK KESSLER has academic degrees in philosophy, law, and library and information studies, and has pursued these and other subjects at Yale, Oxford, and the University of California. He spent fifteen years in the handicraft importing business, until he found the glamor of international travel to be at odds with the joys of married life and of the raising of two small boys. His love affair with books and love/hate relationship with the computer are long-standing. While an importer he fought the automation battles of the 70s and 80s, most often siding with the Luddites against the machines but then reluctantly giving in. He's still suspicious. Currently he works as a networked information consultant, and has just concluded a one-year study in France of the French Minitel and of foreign library applications of the US Internet. He is a member of the American Society for Information Science, the American Library Association, and the California Library Association. His ambition in life still is never to take another airplane trip. His Internet e-mail address is: kessler@well.sf.ca.us

Ways of Connecting to the Internet

by

John S. Quarterman and Smoot Carl-Mitchell,
Texas Internet Consulting

Introduction

Connecting to the Internet should be easy, given all the work in recent years by the IETF working groups on host and router requirements documents, by the User Services working group on FYI documents, and by many people on a plethora of books for users, plus the proliferation of NICs and NOCs. Well, almost. The requirements documents make it more likely that network software and hardware will provide IP, TCP, UDP, and a rudimentary collection of applications. The FYIs and the books help users once they're connected. The NICs and NOCs are very useful to people who know they exist.

But who do you contact to get a connection? What do you need to register, with whom, and how? Is there an Internet, Inc. you can call for service, in the way there's a CompuServe, Inc.? How much does it cost, and who do you pay? What applications do you want, where do you get the software, and how do you set it up? These are some of the questions people new to the Internet ask. We get them all the time in our consulting practice, and we are writing a book about them, *The Internet Connection: A Guide to Connectivity and Configuration*, forthcoming from Addison-Wesley. We'll address some of these questions in later articles in *ConneXions*, but here let's discuss a more basic question, adapting and summarizing the whole chapter it occupies in the book.

Different networks, Different services

What is an Internet connection? The name, the Internet, has become rather vague, being used for everything from direct IP connectivity to dialup UUCP or FidoNet connectivity. The user who pays fifty bucks to a BBS (bulletin board system) that promises "Internet connectivity" may be disappointed to discover that the BBS really only exchanges mail with the Internet a couple times a day, not *Telnet* (remote login), *FTP* (*File Transfer Protocol*), *archie* (index searches of public FTP archives), *WAIS* (*Wide Area Information Servers*, with full document searches) *gopher* (a menu-oriented front end to many other services), *WWW* (*World Wide Web*, a distributed hypertext system and front end to other services), *talk* (interactive one-to-one communications), *IRC* (*Internet Relay Chat*, interactive many-to-many communications), *finger* (check a person's personal description), *netfind* (find a person), and the long list of other services that true Internet connectivity provides. This is why it matters which network you're connecting to: different networks provide different services, as one of us has previously discussed in *ConneXions* ("Network Nomenclature," September 1990). Let's adopt more concrete terminology, for this article at least.

If you can FTP to `is.internic.net`, `ftp.psi.com`, or `ftp.uu.net`, you're *on the Internet*; otherwise, you're not. Somewhat more formally, *an internet* is a collection of networks that communicate with each other using the same internetwork protocol, and *the Internet* is the largest internet using the *Internet Protocol* (IP). The Internet is actually multi-protocol these days, but we don't have space here to discuss OSI, AppleTalk, NetWare, DECnet, SNA, or even IP over IP. (We won't even define these names, except to say they are other network protocols, capable, no doubt, but different.) In addition, new users are interested in what they can do with the network, that is, what applications they can use, and all the applications we recommend here run over TCP/IP or UDP/IP, and many of them only over those protocols.

If you can send mail to someone on the Internet through UUCP, FidoNet, NJE, or some other means, but you can't FTP, you're not on the Internet. However, you are *in the Matrix*. *The Matrix* is all computer networks worldwide that exchange electronic mail. It includes the UUCP mail network, *FidoNet*, *BITNET*, commercial mail systems such as *MCI Mail*, and conferencing systems like *CompuServe*, *GENie*, *BIX*, *the WELL*, *the World*, and many others. Another common service in the Matrix is *USENET news*, which is an asynchronous many-to-many communication service that distributes articles to machines by a flooding algorithm, rather than messages to personal mailboxes, as mail does. *USENET* isn't really a network; it's just the set of all computers that get *USENET* news. Since any host that gets *USENET* news almost certainly exchanges mail with other systems and is thus in the Matrix, we don't say a lot about news here.

Many of the networks and systems in the Matrix provide other services, but the least common denominator is electronic mail. By this we mean mail to and from other systems, not just internal mail. Many thousands of stand-alone BBS systems are not part of the Matrix, since they don't exchange mail with any other systems; some corporate enterprise networks are still not part of the Matrix for the same reason; and the last we heard *Prodigy* was still not in the Matrix.

The mail end of the spectrum

There is a spectrum of connectivity between the least common denominator, mail, and full Internet connectivity. Mail-only systems provide the least service, but are also usually the least expensive. *UUCP* is the name of a mostly dialup mostly UNIX network that mostly provides only mail and often *USENET* news. It is named after the UUCP (UNIX to UNIX CoPy) protocol that it uses. *FidoNet* is a mostly dialup, mostly DOS network that provides mostly mail and often echomail. *Echomail* is a service much like *USENET* news, and that often carries articles from *USENET* newsgroups (discussion topics). *UUCP* and *FidoNet* reach countries, cities, organizations, and people that other networks do not, because these two networks are among the least expensive, yet most capable for what they cost. Traditionally, the way you connect to either network is to get someone to agree to let you connect to their machine, after which you pay the telephone bill for the immediate connection. Now there are also commercial services, such as *UUNET*, *PSI*, *EUnet*, *JUNET*, and many others, that will provide you *UUCP* connectivity for a fee.

Commercial conferencing systems, such as *CompuServe*, *BIX*, *GENie*, and others, often exchange mail with other systems, and are thus part of the Matrix. Many of these systems actually have IP connections to the Internet, but deliberately limit traffic over such connections to mail only. Mail to these systems is thus often faster than mail by *UUCP* or *FidoNet*, but it's still just mail. Yet you may find a commercial conferencing system more to your advantage than a mail network such as *UUCP* or *FidoNet*, because such conferencing systems often provide a wide range of other services, such as communication forums, databases, stock quotes, airline scheduling, etc. These other services often cost more, on top of an hourly connect charge, but they may be worth the cost to you.

The IP end of the spectrum

For actual Internet access, there are at least two major variables:

- Are you using a dialup or direct connection?
- Is your own computer or network connected to the Internet?

Ways of Connecting to the Internet (*continued*)

These are almost orthogonal criteria. You can dial up a host (let's call it a *login host*) that is on the network with nothing but a good modem, and that host can have full Internet connectivity with a full range of services, but your host won't. Or you can use *SLIP* (*Serial Line Internet Protocol*) or *PPP* (*Point-to-Point Protocol*) to support IP over a dialup modem connection (let's call this *dialup IP*), thus providing your own host with Internet connectivity.

Using a login host lets the administrator of that machine handle setting up applications, doing backups, and other administrativia. You can FTP files to the login host, but if you want a file on your own machine, you must use some other means, such as *Kermit* or *Zmodem*, to get the file over the dialup modem hop to your machine. You can probably use WAIS, archie, Gopher, WWW, etc., but most likely only through interfaces provided by the login host.

Using dialup IP requires you to get and install IP, TCP, and UDP software on your machine, plus SLIP or PPP, plus various application packages. Fortunately, packages with the appropriate software are available both in free versions and in commercial vendor-supported versions. The advantage is that you can FTP directly to your own machine, and you can use directly as many of Telnet, Gopher, etc., as you are willing to set up on your machine.

Full Internet connectivity

Full Internet connectivity usually involves some sort of dedicated IP connection. Here there is another spectrum, between frequent SLIP or PPP connections with a high speed (14.4 kilobit/sec) modem and a T3 (45 megabit/sec) 24 hour dedicated pipe. Given the recent dramatic reduction in costs of high speed modems, and the spread of newer potentially intermittent services such as ISDN, Frame Relay, or SMDS, these distinctions have become fuzzy. Basically, the more you're willing to pay, the faster and more dedicated the IP pipe to the Internet you can get.

There is a social aspect to full Internet connectivity, as well. Unlike with conferencing systems or traditional passive consumer communications media such as newspapers or television, the Internet is a participatory medium. You can not only use services provided by others, you can provide your own services for others to use.

Summary

Let's summarize this spectrum of network access with a table:

Type:	Matrix Access		Internet Access		
	mailnet	conf	Login Host	Dialup IP	Full
mail:	yes	yes	yes	yes	yes
news:	yes	maybe	yes	yes	yes
FTP:	no	no	yes	yes	yes
Interactive:	no	no	yes	yes	yes
IP to:	gateway	gateway	login host	your machine	your machine
dialup:	yes	yes	yes	yes	yes, or dedicated
speed:	modem	modem	modem	modem	modem and up
cost:	*	*, †	*	*	monthly

mailnet mail and news only, e.g., UUCP or FidoNet

conf conferencing system with mail and news access

* monthly + connect time charges

† per message charges

We're not interested in systems that are not part of the Matrix, since they don't even have mail access. We omit such systems from the table, so all types of systems described in the table have mail access. Most of the systems described in the table can provide USENET news or something like it if you want it, although some conferencing systems do not provide conferencing outside of their own machine.

The boundary where the Internet begins within the Matrix is determined by IP connectivity, either to a computer you have dialed up, or to your own machine. This is seen in the table in the rows about FTP and interactive access. Mail networks and conferencing systems don't have either.

If a conferencing system does permit FTP and other interactive IP services, we simply count it as a login host. A login host is on the Internet, and thus you have Internet access, but your own machine is not. If you connect your own machine for dialup IP, your machine is then on the Internet, but in a limited manner. You can consume Internet services but it is difficult for you to provide them, since your machine is not connected all the time. Full Internet connectivity may be accomplished with either dialup or dedicated links, and at speeds ranging up to very fast.

In our consulting practice, we find about 80% of those who say they want to be "on the Internet" really just want mail. For them, we recommend UUCP, FidoNet, or a commercial conferencing system, any of which can put them in the Matrix without the cost of being on the Internet. For those who really do want to be on the Internet, so they can use FTP and other Internet services, we often recommend a login host. Others want more direct Internet access, and may use dialup IP, or may get full Internet connectivity.

The value of the Internet community

Which access method people choose usually depends on how much Internet services are worth to them, and how much they are willing to pay. These are not easy questions to quantify. The real value of the Internet is not any one service: it is the participatory community. This is a difficult concept to explain to someone who has not experienced it. Fortunately, a spectrum of connectivity is available that permits new users to gain experience at minimal cost, and increase connectivity later.

SMOOT CARL-MITCHELL and **JOHN S. QUARTERMAN** are partners in Texas Internet Consulting (TIC), which consults in networks and open system with particular emphasis on TCP/IP networks, UNIX systems and standards. They also write articles and books and give tutorials. They have just published an in-depth examination of *Practical Internetworking with TCP/IP and UNIX*, and are finishing *The Internet Connection: A Guide to Connectivity and Configuration*, both from Addison-Wesley, 1993. Quarterman is also a co-author of the books *The Design and Implementation of the 4.3BSD UNIX Operating System*, 1989, and *UNIX, POSIX, and Open Systems: The Open Standards Puzzle*, 1993, both from Addison-Wesley, and is the author of *The Matrix: Computer Networks and Conferencing Systems Worldwide*, Digital Press, 1990. Quarterman and Carl-Mitchell are the editor and managing editor of *Matrix News*, a monthly online and paper newsletter about contextual issues crossing network, geographic, and political boundaries, published by Matrix Information and Directory Services, Inc., of Austin, which also publishes *Matrix Maps Quarterly*.

This article is adapted from a forthcoming book, *The Internet Connection: A Guide to Connectivity and Configuration*, Addison-Wesley publishers, ISBN 0-201-54237-4, 1993. Used with permission.

Establishing Workgroups in a Multiprotocol Environment

by Bob Meindl, Booz-Allen and Hamilton

Abstract

The move to downsize to client-server architectures coupled with the increasing need for formerly isolated LANs to share information has forced many sites into running multiple protocol families over a single corporate backbone. This move has been enabled by the current generation of multiprotocol routers which allow several protocol stacks (along with their attendant routing protocol or choice of protocols) to operate simultaneously within a single box. Some sites faced with this problem have chosen to bridge networks together or to encapsulate data within a single protocol on the backbone. This article examines the establishment of *workgroups* when several protocols are being routed within a single backbone structure. Until an integrated routing protocol can support the majority of today's popular protocols, establishing workgroups requires that topological as well as addressing issues be resolved individually for each protocol using the "Ships in the Night" [1] approach. Reasons for establishing workgroups are usually traffic based, security based, or both. A general outline for attacking this problem at the network (i.e., router) layer is presented, followed by specific discussions in the context of a sample network for TCP/IP, DECnet Phase IV, Novell SPX/IPX and AppleTalk.

Why establish workgroups?

There are several reasons to segregate workgroups in today's inter-networking environment. Two of the predominant reasons are traffic segregation and security. Traffic segregation addresses issues such as efficient bandwidth utilization on an over-taxed network, efficient routing according to some metric (e.g., cost, number of hops), user quantity, and absorption of new networks into the corporate internet. Any network administrator who has seen network performance plummet as the number of attached devices increases is familiar with the technique of splitting the network into multiple subnets through the use of bridges and/or routers. This is done due to the limitations of current technology (i.e., Token Ring, Ethernet, LocalTalk) in handling large quantities of interactions between many participating stations.

As these corporate networks have grown, a large percentage have evolved to use of a corporate backbone with attached subnets. The desire to minimize backbone traffic necessitates the identification of groups which utilize dedicated hosts and/or servers for day-to-day activities. These groups may use corporate-wide services such as electronic mail and therefore need access to the corporate network, but they routinely work with their own servers which other groups don't need to access.

Another reason to establish workgroups is the need for security. Security in the context of this article refers to network configuration-based restrictions on data access, not specialized security applications. A familiar example might be the company which has divided their networks along functional lines such as marketing, engineering, personnel and accounting. Engineering personnel generally don't have a need to access the personnel database (although they may like to), and the establishment of restrictions on workgroup boundaries can provide a first line of defense in restricting access to data.

The bottom line is that for whatever reason an organization decides to isolate traffic, a methodical approach to establishing workgroups during the network design or network re-design process can ensure that the network structure is well thought out, understood by network administrators, scalable, and maintainable.

Attacking the problem

It is initially assumed that the network in question has grown to the point where there are separate subnets connected to a corporate-wide backbone. The subnets could consist of a mixture of various data-link technologies; Ethernet, Token Ring, LocalTalk, FDDI, ATM etc. The backbone could also consist of any one of these. This article discusses an architecture consisting of an FDDI backbone with Ethernet subnets but since the following discussion pertains to routing (i.e., the network layer) there is no loss in generality in using these technologies as a reference.

Figure 1 depicts a generic network consisting of a backbone interconnected via attached routers, with each router providing an attachment to one or more specific subnets. In the ideal case, both the physical and logical network architectures are designed around an organizational structure but this is often not the case due to 1) constraints imposed by the physical building(s) in which the networks are located and 2) moves, adds, and changes in a rapidly changing organizational structure. The more prevalent case is the one in which the overall network architecture and media are already chosen due to preliminary bandwidth considerations and growth requirements, and the subnet structure is already chosen (if not in place) based upon building layout, signal closet location, etc.

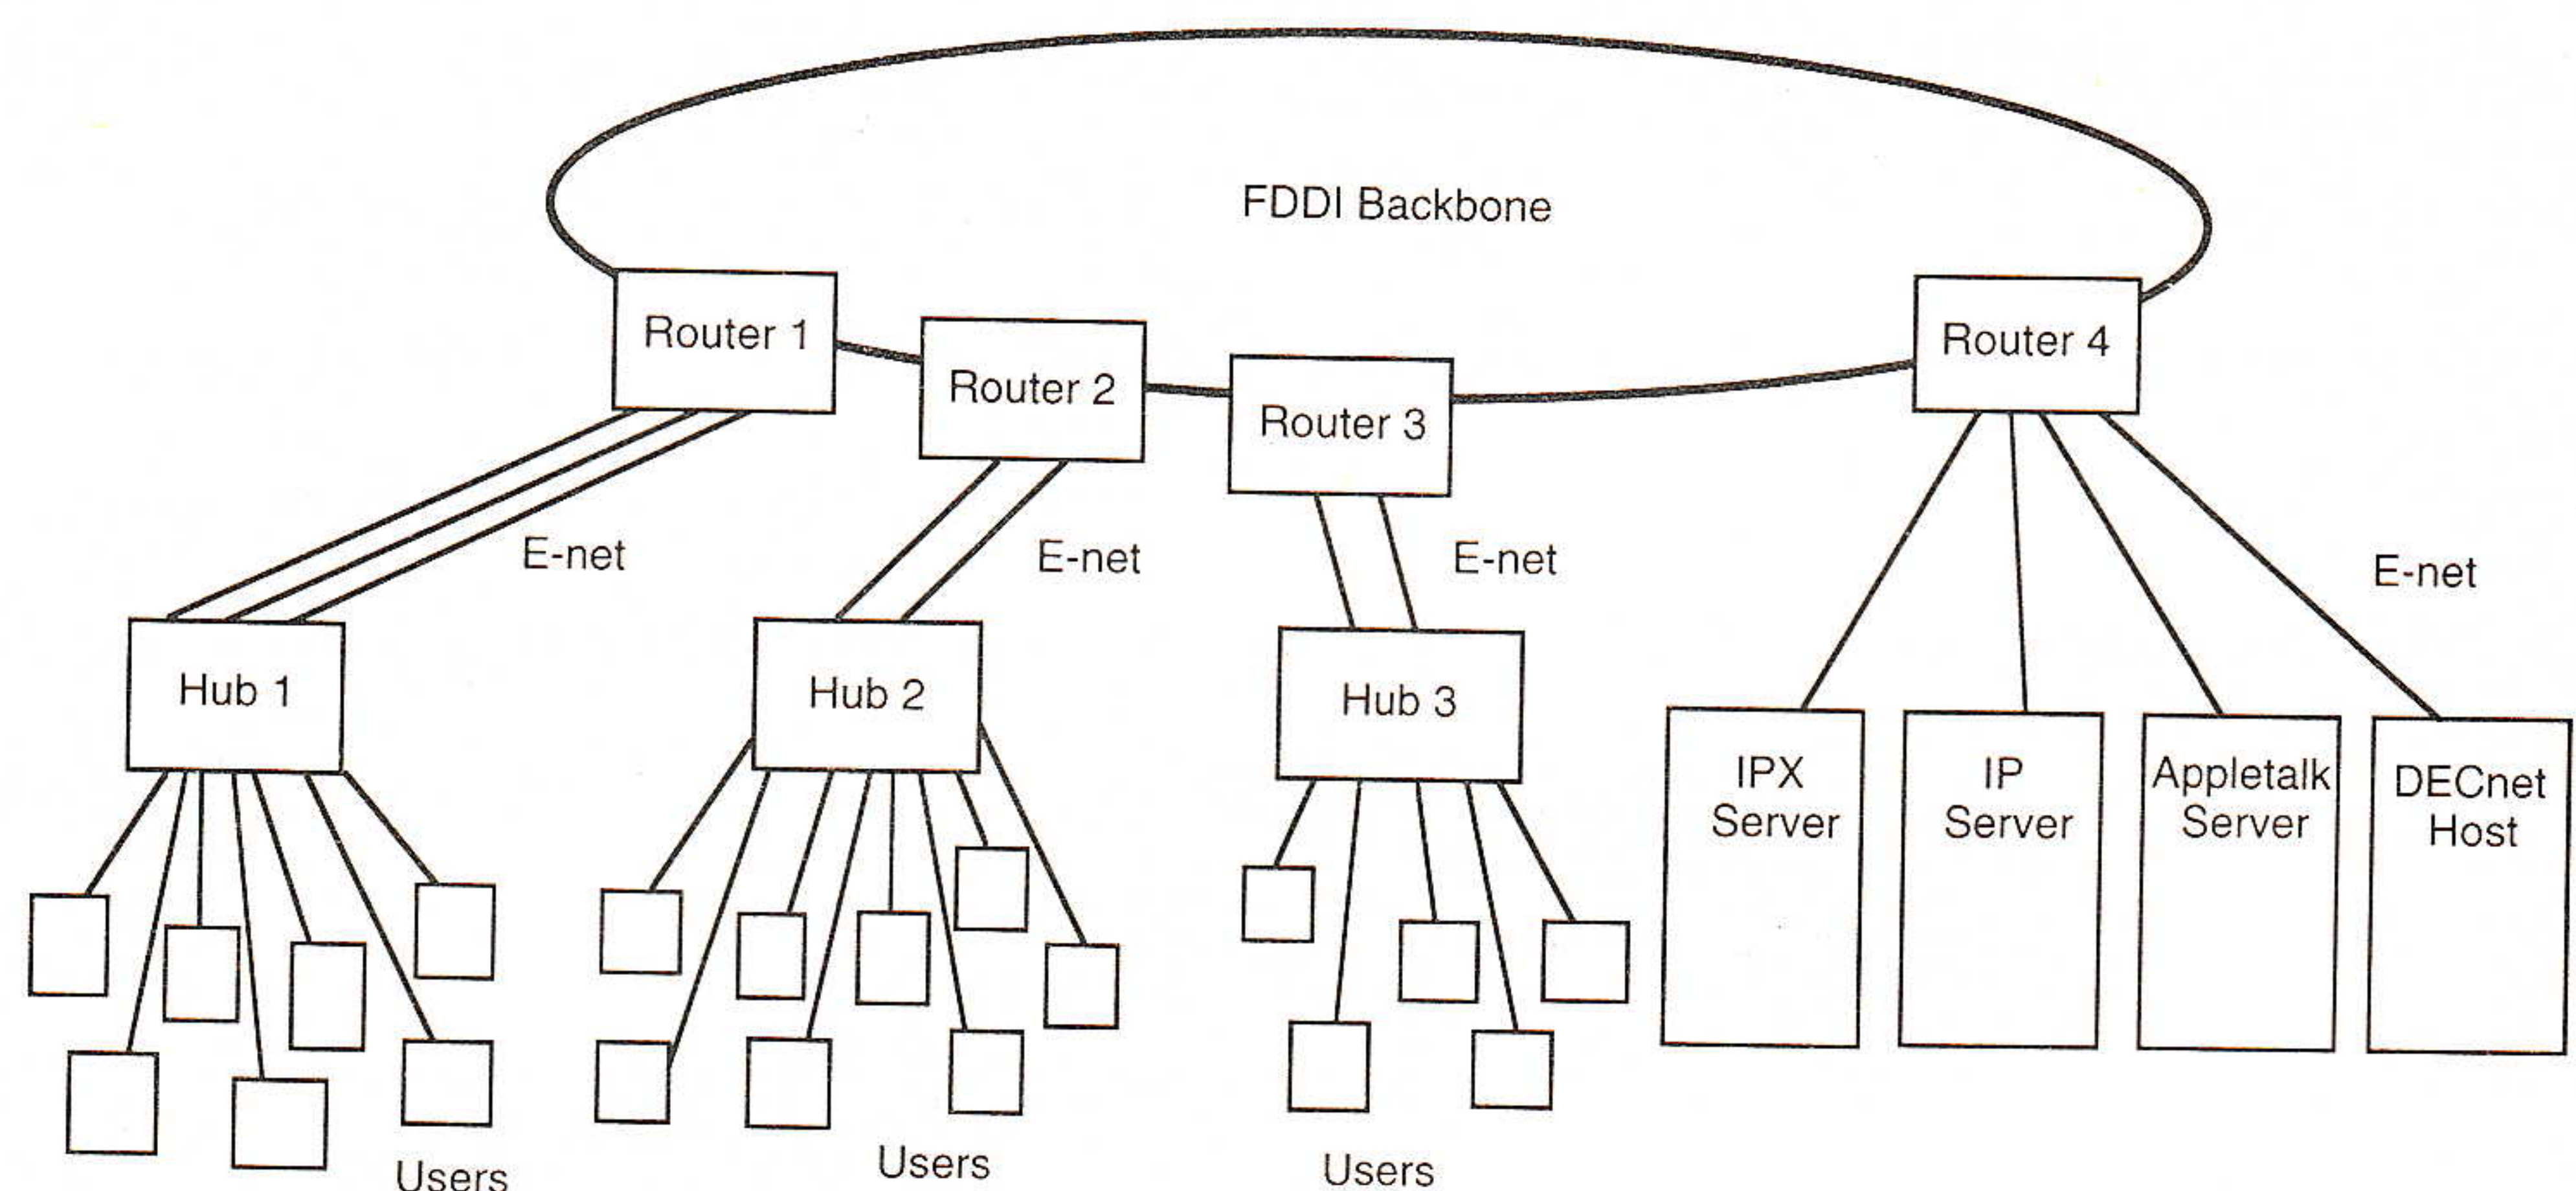


Figure 1: Sample Network Physical Topology

User stations or servers on any subnet may contain a mixture of upper layer protocols on a physical media such as Ethernet. Also, some stations may contain dual stacks, where dual stacking refers to the practice of a single station running two upper layer protocol families simultaneously (e.g., TCP/IP and SPX/IPX). End stations using different upper layer protocols may share the media, but do not "see" each other nor can they directly communicate with one another. Another case which needs to be taken into account is the one in which servers/hosts providing a service to users are not located on the same subnet as the user. A mixture of local servers and remote (where remote refers to a server that is not on the local subnet) servers may also be utilized by users on a given subnet. A further extrapolation of this case is a server located on a local subnet which is accessed by the users on that subnet as well as users throughout the organization. Each of these cases needs to be taken into account in establishing workgroups.

Logical network architecture

Establishing Workgroups (continued)

Another important consideration in the establishment of workgroups is the network's logical architecture. The logical architecture refers to the network topology as seen by the routing protocol(s) and data packets as opposed to the physical topology which represents the physical interconnection of network devices. An example of this would be an Ethernet subnet connected through an intelligent wiring hub. While the physical view of this would be star-wired, the logical view is of all stations on the subnet connected to a common Ethernet subnet. A simple example of this is shown in Figure 2, which depicts the logical view of the network illustrated in Figure 1.

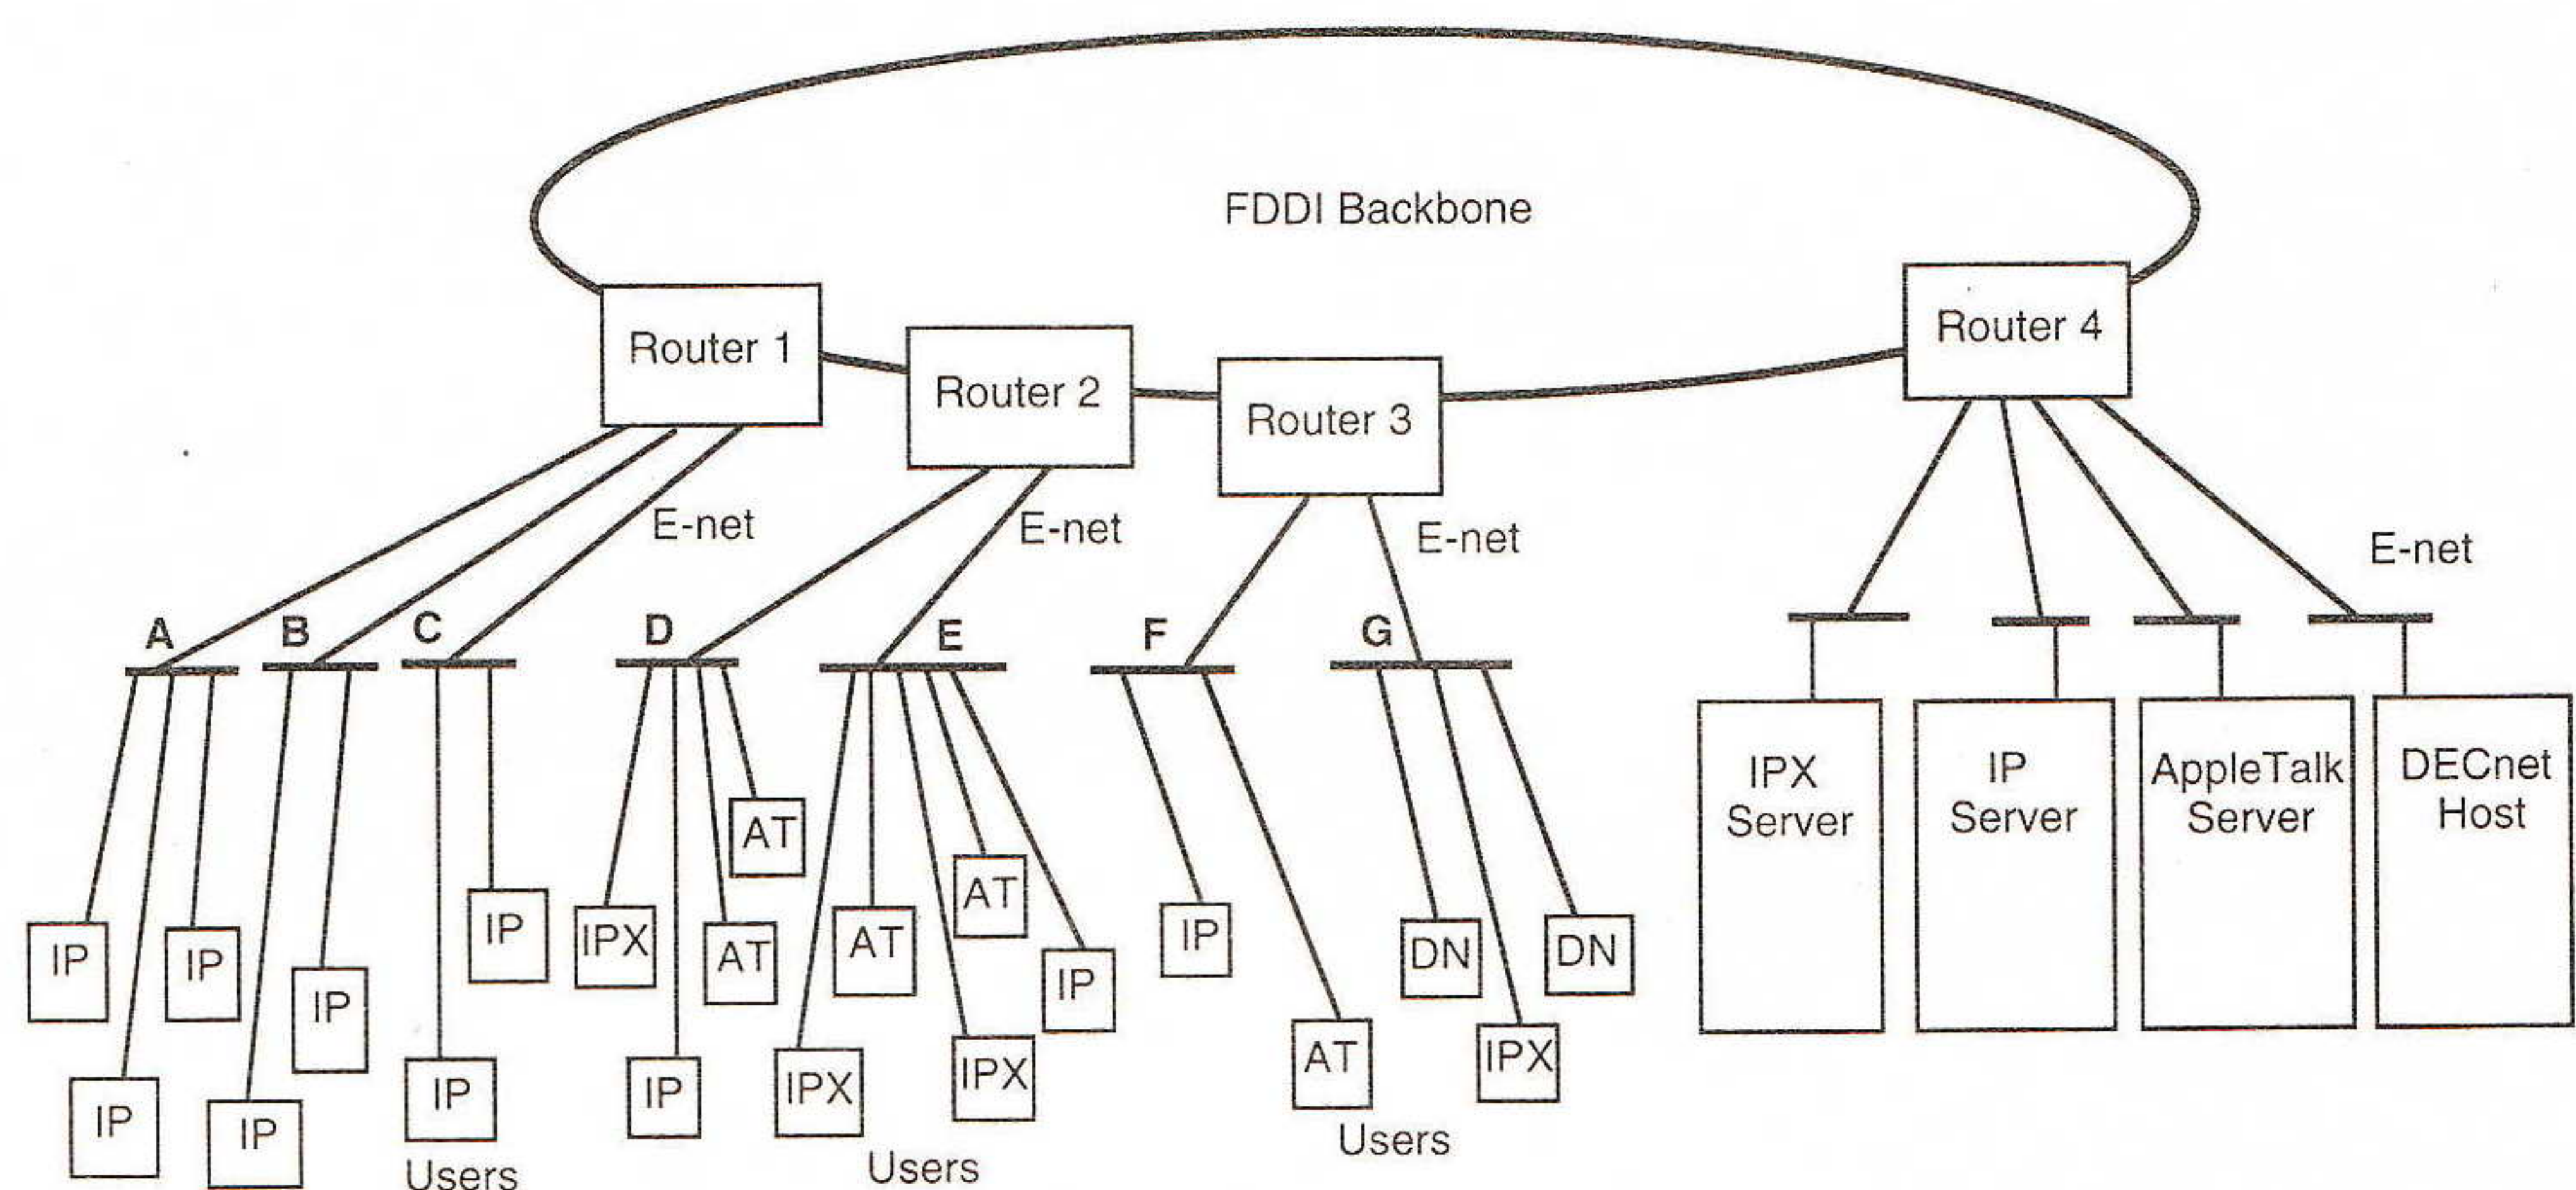


Figure 2: Sample Network Logical Topology

Ships In the Night

Most of today's common protocol families use their own routing protocols to determine the best paths through a network and to enable inter-router communication (e.g., TCP/IP with RIP, OSPF, or Integrated IS-IS; SPX/IPX with Novell RIP or NLSP, etc.). One method in which a multi-protocol environment could be established would be to use a separate router designed for each routing protocol. Most of today's multi-protocol routers, however, allow routing protocols to run side-by-side but independently from one another in a single router. Routing tables, inter-router communication, etc. are maintained separately for each protocol enabled within the router. This approach is termed *Ships in the Night* [1] routing. Figure 3 depicts the appearance of the network from the perspective of the IP routing protocol. This principle would then be individually applied to each routing protocol in use. The alternative to this approach when a single, multi-protocol router is used is "integrated" routing in which one routing protocol supports multiple protocol families. An example of this is Integrated IS-IS, which currently supports CLNP and IP routing [1].

General methodology

In many cases, the network architecture either already exists or has been chosen before anyone gets to thinking about workgroups. Also, workgroups may already exist in the form of disjoint LANs which are just being brought into the enterprise network. In either case, a specific methodology should be employed in establishing workgroups within the framework of a multiprotocol environment. In reality, the methodology identified below is not carried out in the order presented as it is an iterative process that may need to be repeated several times to refine the workgroup design.

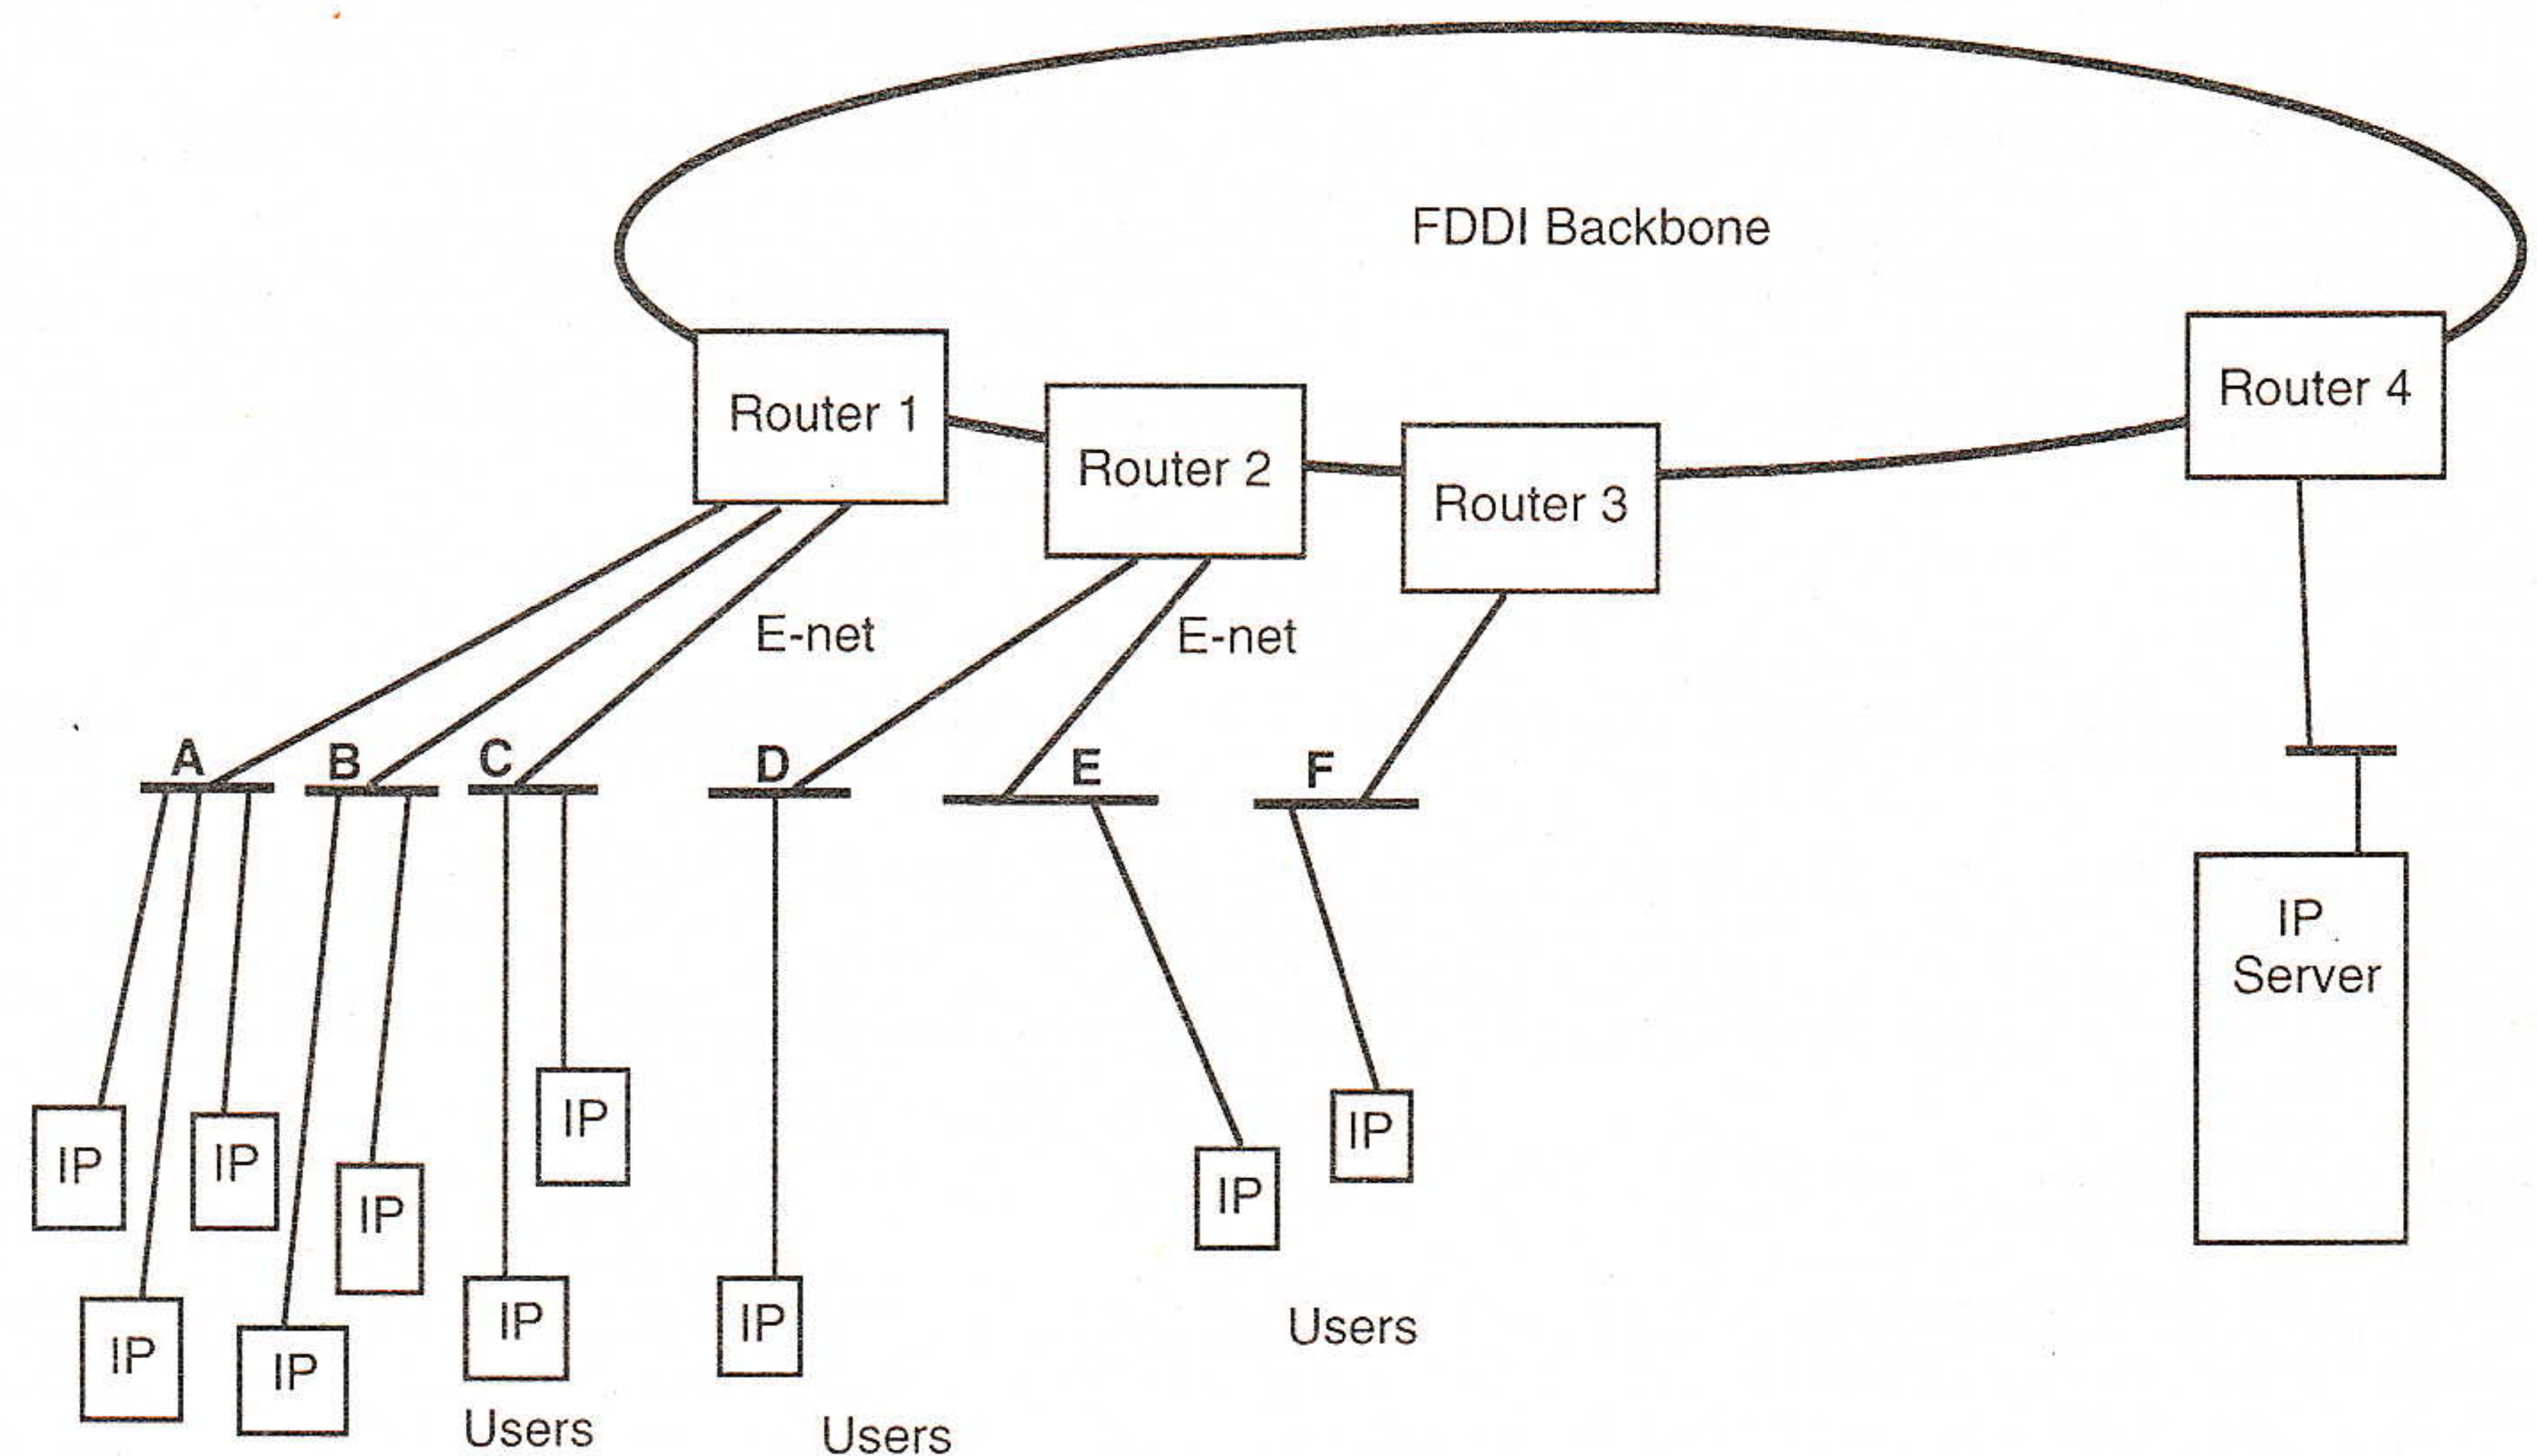


Figure 3: IP View of Logical Topology

1. Understand the connectivity of the current/proposed network structure including the physical and logical architecture. If it doesn't already exist, sit down and draw a block diagram of the network showing physical connections between routers, hubs, groups of users, servers and hosts. This will assist you in understanding the logical architecture, which is probably the most important network view in establishing workgroups. Next, draw an equivalent diagram of the logical architecture which shows which devices communicate with one another, and which network devices communicate to groups of users, servers and hosts.

2. Determine current and planned protocols to be used on the network. This is important because of the subtle and not so subtle differences in routing protocols which are used within each protocol family. There are several approaches to doing this, where the more common approach is to use a combination of the methods discussed below. The approach taken is dependent upon the size of the network as well as the size of the user community.

- *User Survey:* In smaller groups it may be feasible to survey all users as to their network connectivity needs. This would include what protocols are being run and what servers/hosts are being accessed. Often this approach is not productive because many users have no idea of how they're connected to their server(s). An alternate to this method is to have the local network gurus perform the survey by inventorying machines. This approach is costly and time consuming, however.
- *Host/Server Survey:* The network administration group can survey the hosts and servers to determine all protocols running on the machines. If a stack isn't running on a host and/or server then a user won't be able to make use of it even if the stack is installed on a user machine (unless peer-to-peer networking is used). This approach is less time consuming and probably more accurate than the user survey method.
- *Network Analysis:* If a current enterprise network is being upgraded or a new group is being assimilated into the enterprise, then a good network analyzer attached to the network(s) can tell you what protocols are being run during a given time interval.

continued on next page

Establishing Workgroups (*continued*)

If all protocols in use are to be identified, it is important to run the analysis over several days and maybe even a week to capture all protocols. The downside of this is that in a large, complicated network, little-used but important protocols could be missed and therefore not considered in the analysis.

3. If not already determined in step 2, associate the protocol or protocols in use with specific hosts and servers.

4. Further refine the user groups by associating them with specific servers, hosts and specific services provided by those devices. Determine if the group is serviced by one subnet, one router, or is dispersed throughout the enterprise. Begin to correlate the logical and physical connectivity by analyzing physical locations of user groups and determining the router their subnet is attached to. Also, separate enterprise-wide services (e.g., e-mail) from group-specific services and applications.

5. Using the list of protocols found earlier, study the routing protocol in use for each protocol family and the layer three packet header format in detail. The methods used by routing protocols to pass information, and the packet header format will predetermine most of the useful "filters" which can be used to restrict traffic. This data can be found in various places such as Internet RFCs, vendor published material, and textbooks. A partial listing of references is provided at the end of the article.

6. Understand the capabilities of your router(s) with respect to each protocol. This includes functionality as well as performance when filtering is enabled. During network design or a major upgrade, it is often useful to prototype routers that are under evaluation with filtering enabled, in a similar configuration to that in which the router will be used. Methods for configuring and performing these tests are the subject of an upcoming paper.

7. Filters should be placed as close to the workgroup as possible. This will mitigate any performance degradation which could occur with excessive filtering and will minimize unnecessary traffic on the backbone.

Example

To construct good filters, it is important to understand the network-layer header, the information it carries, and the method used by the routing protocol to determine routes based on that information. The references at the end of the article explain the header format for the four protocols to be discussed in this example. Any of these fields can be used as a basis for filtering however different routers may have some of these available as predefined filters to ease configuration. Some routers also provide the user with the ability to configure custom filters based on arbitrary bit patterns in the header. The design and configuration of these filters is heavily dependent upon the specific router in use. It is also important to note that different routers may experience different levels of performance degradation (from trivial to major) when predefined or user-defined filters are implemented.

The routing protocol in use can also influence the choice of filters in that the routing protocol may already provide some degree of traffic segregation. This is particularly important when designing filters for AppleTalk networks, which dynamically assign node IDs and network IDs [6].

This dynamic assignment necessitates that actual routing protocol transactions rather than network or node IDs be used as a basis for filtering. Understanding the routing protocol in use for each protocol family can help to make more intelligent decisions with respect to filter design.

The example network shown in Figure 1 consists of four routers interconnected via an FDDI ring, with one router dedicated to host, server and wide area connections and the other three to subnet connections. Hubs are present between the user and router to minimize router port requirements, provide structured wiring in the physical network, and provide centralized management points. Also, additional routers could be present between the backbone-attached routers and the hubs but these are not in the example for simplicity. There is no loss in generality in the methods used below if this is the case.

The association of user groups to specific hosts, servers, and services can be more difficult because of the need to get detailed data on a per user basis. This can be mitigated if there is a departmental network representative from which this data can be gathered. It is also easier to construct this list on a group boundary versus an individual user boundary, with the understanding that users utilizing services and devices outside of the group norm may lose connectivity upon initial implementation of the workgroup scheme, and modifications will need to be made to accommodate this. Several example workgroups are listed below and shown in Figure 4 on the following page.


- A: Group on one segment with local server
- B: Group on several segments served by one router & local server
- C: Group on several segments across several routers with local router on one of the segments in the group
- D: Servers/hosts centrally located across backbone

Case A In Case A, general workgroup traffic would remain on the local segment and the only filters necessary for IP would be security-based to exclude non-workgroup nodes from accessing the local server. This filter could be set-up to filter incoming traffic based on network number to exclude all outside networks. Some communication is probably desired for enterprise-wide services hosted on the server connected to router 4 and in this case the filter could exclude all but subnet A and subnet I.

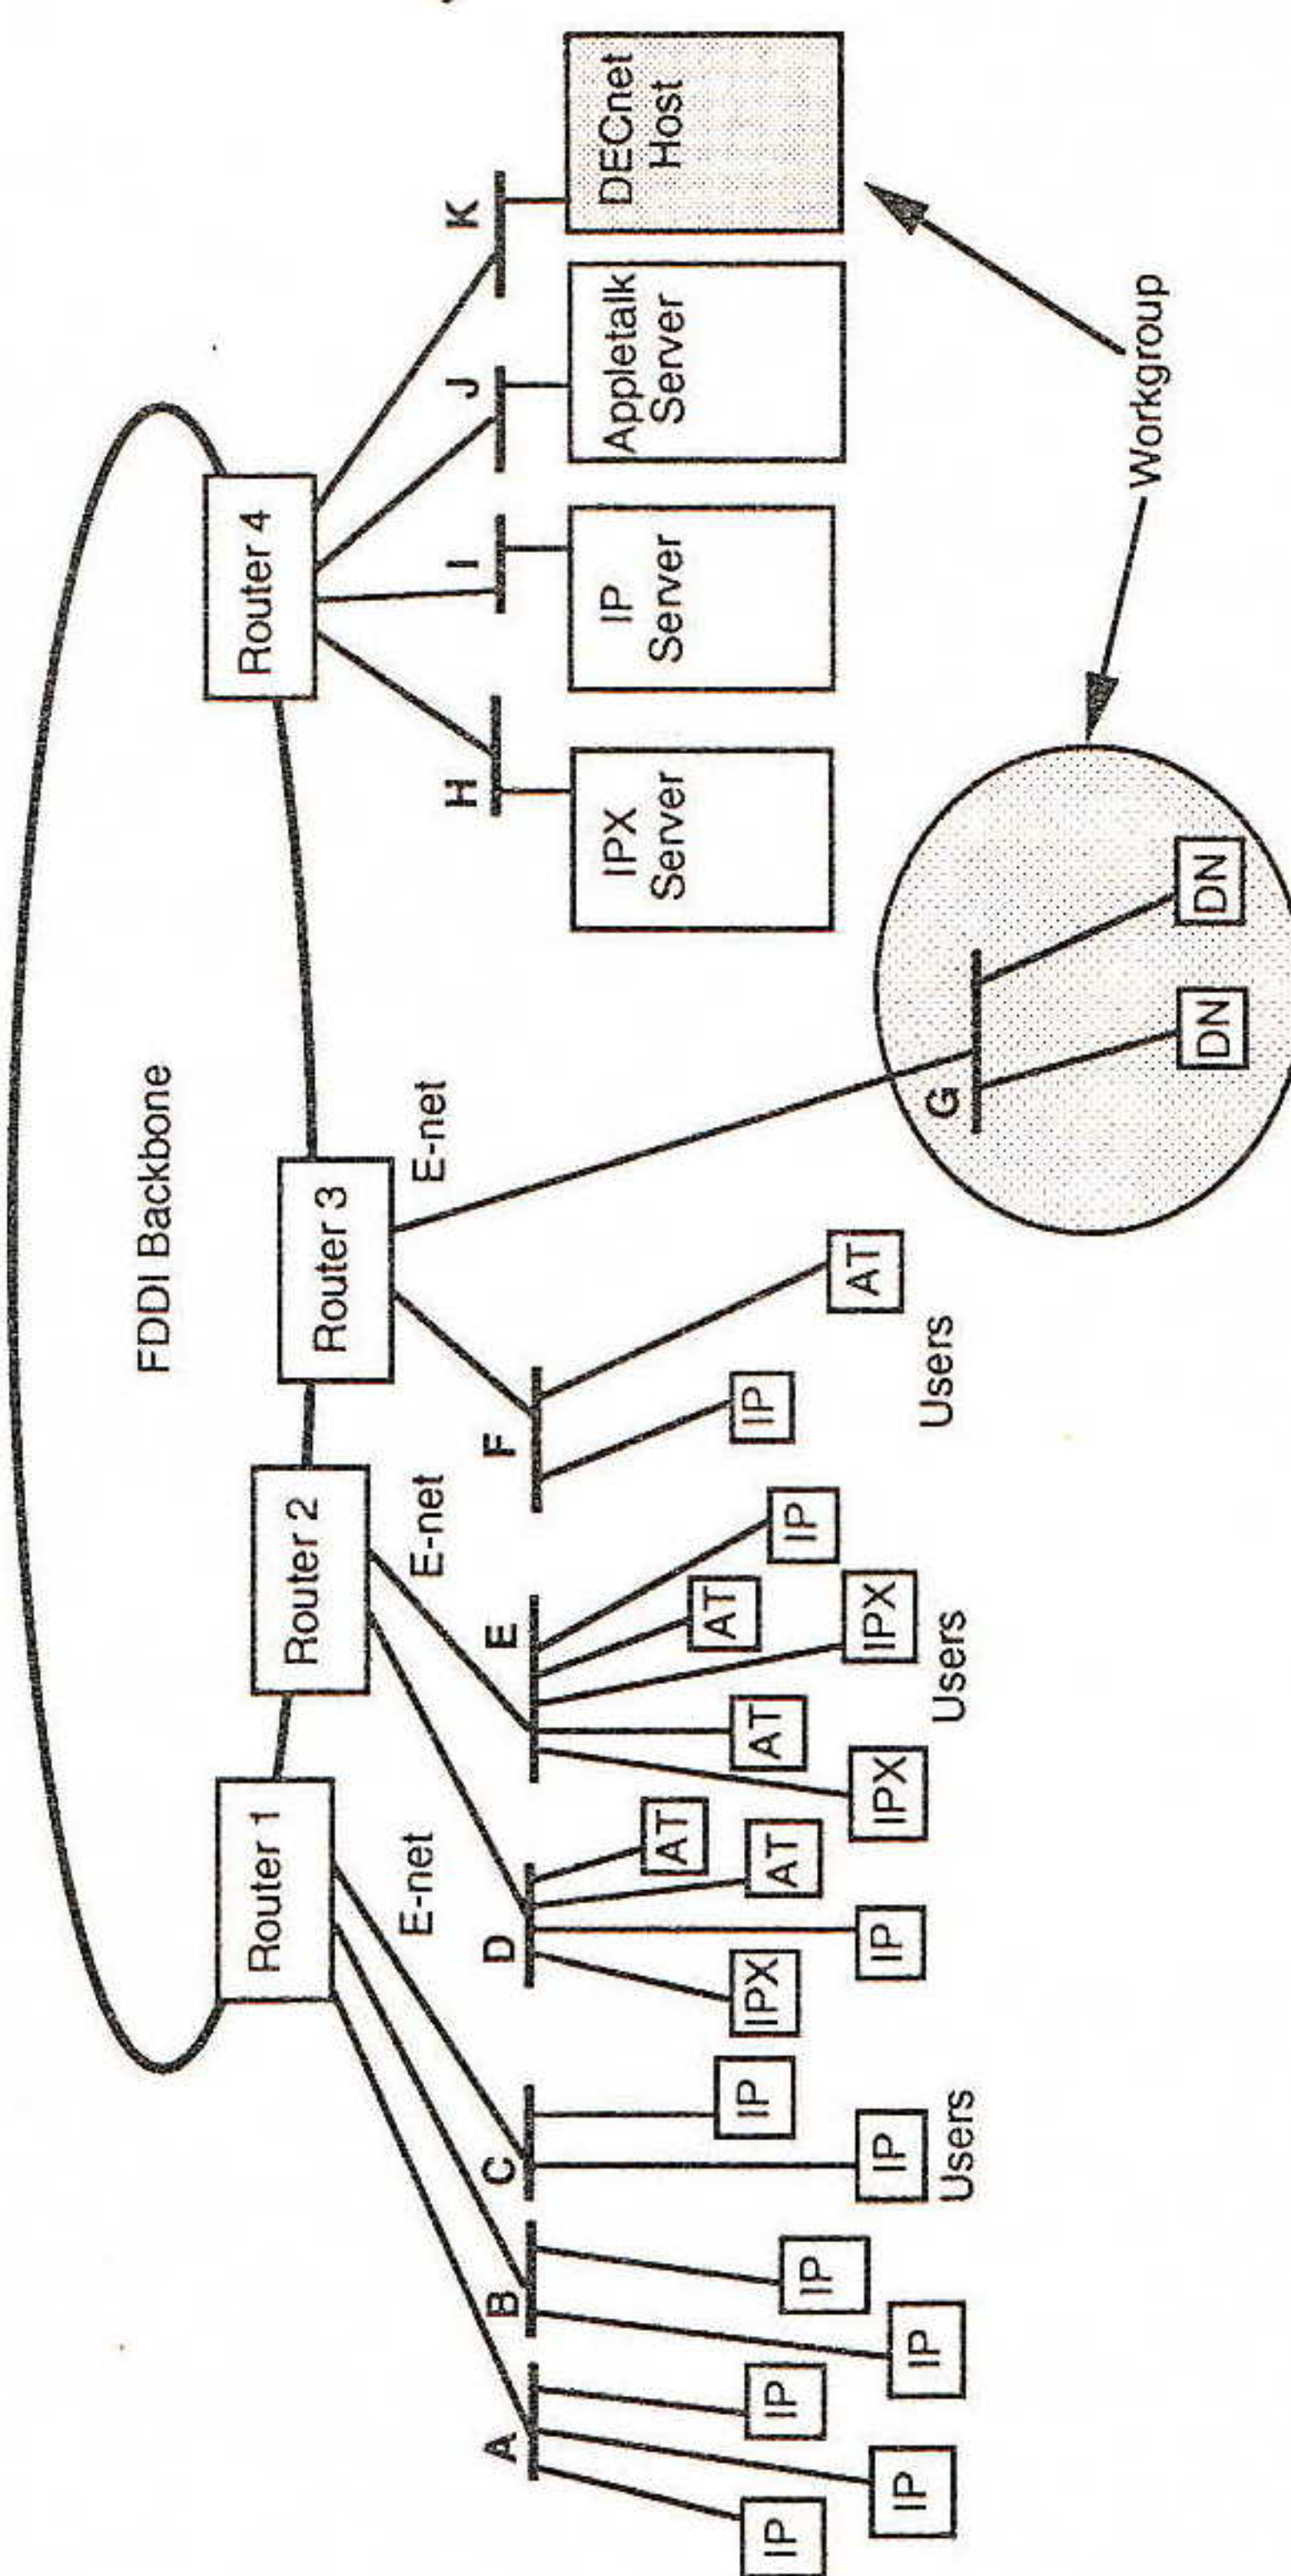
A similar methodology could be followed for DECnet Areas and IPX Network IDs. AppleTalk filtering in this scenario requires knowledge of the zone set-up unless there is a one-to-one correspondence between the physical Ethernet segment and network numbers, in which case the above scheme could be followed [6].

Case B It is also fairly easy to configure the workgroup depicted in Case B from using the backbone since the majority of traffic will remain local to the router. This is a similar situation to Case A except that the router port servicing subnet D would only allow external communication with subnet E and vice-versa. Each of these ports is then configured to allow communication with Subnet H for enterprise-wide service. This case is the one in which good subnet numbering can play a role via the use of address ranges. Some routers allow you to specify a parameter range to filter on, where the parameter can be source or destination address, socket number, host ID, etc.

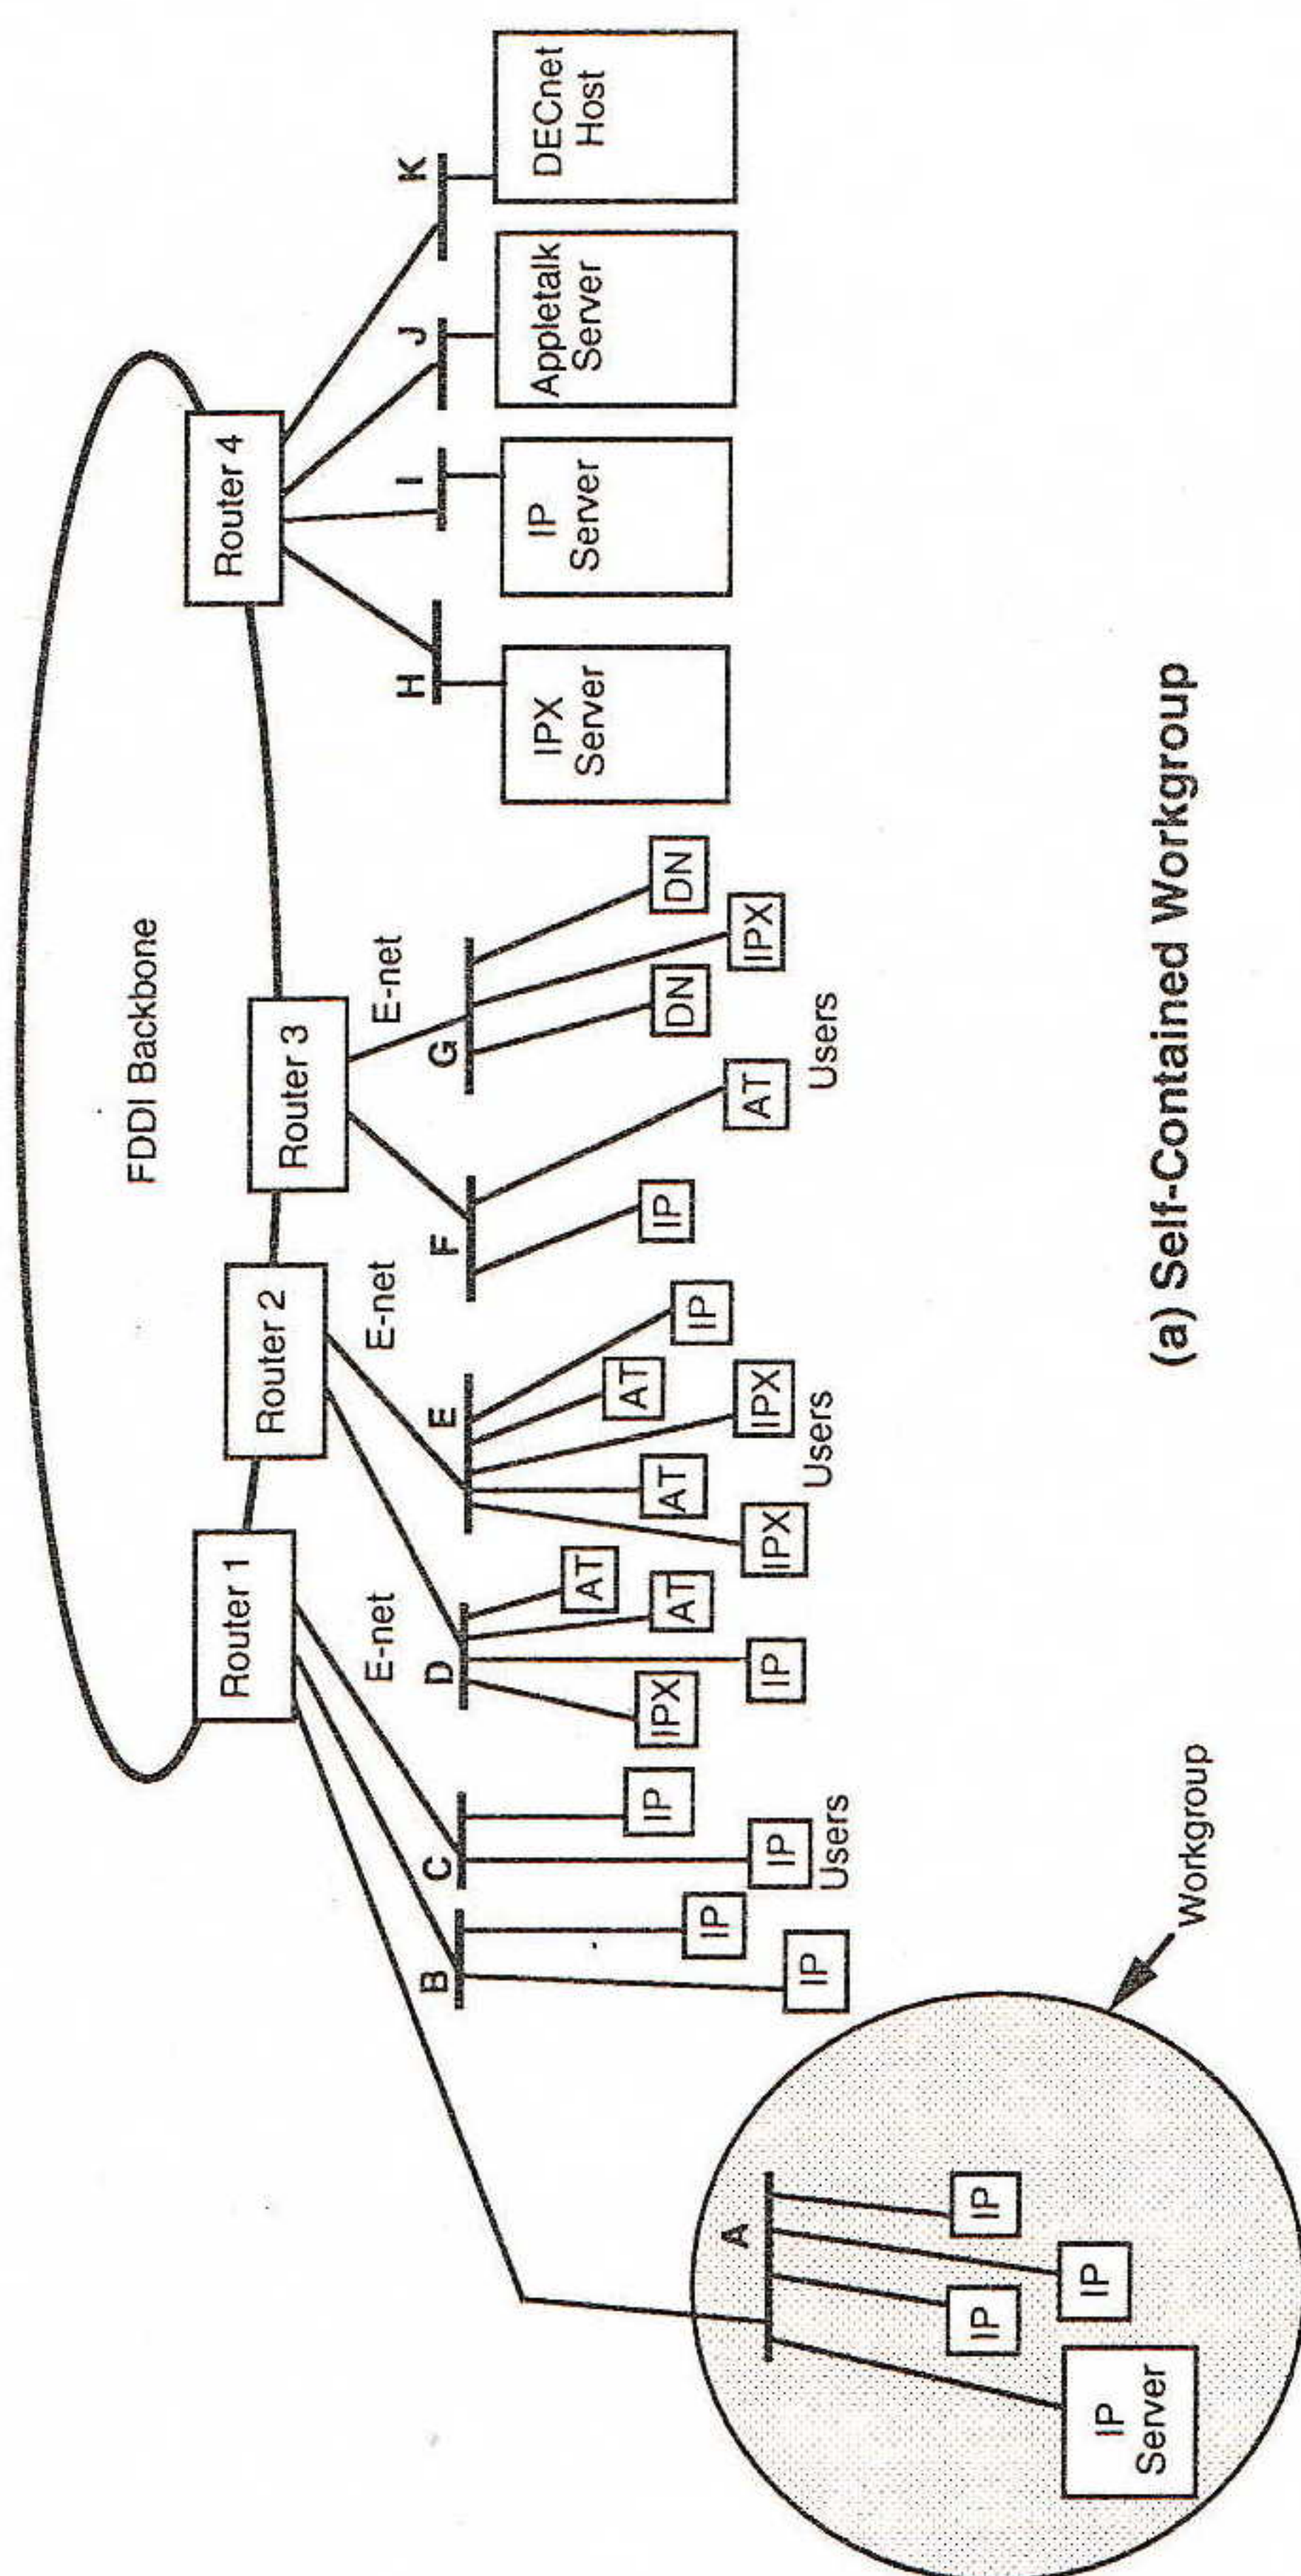
Establishing Workgroups (continued)




(a) Self-Contained Workgroup



(b) Single Router Workgroup with Local Server



(c) Multiple Router Workgroup with Local Server



(d) Workgroup with Centrally Located Server

If all the subnets requiring access to a specific server fall in a range, then you can use one filter to specify that any packets with "parameter" set outside the range are excluded, instead of using multiple filters. This technique can also be used if you have a nearly continuous range by setting up a filter to take the required action based upon the range and then constructing additional filters for the exceptions.

Case C In Case C it is more difficult to keep traffic off the backbone since the workgroup is spread across more than one router. Also, the ports on routers 2 and 3 need to take into account the IP user on subnet B which is not part of the workgroup. This can be done by creating a second filter to drop packets with specific host IDs. In this case, influencing the location of user connection into the network and user addressing can play a role in minimizing the number of required filters. Filtering can then be performed on user generated packets to exclude all packets except those from specific hosts. If there are multiple workgroups present, then preassigning a range of node IDs to each workgroup allows you to filter based on ranges in much the same manner as the previous case, with the difference being the parameter filtered (node ID versus network number). This approach is only practical for IP as the other protocols make use of the preset MAC-layer address for the network-layer node ID [5-7].

Case D It is not practical to keep the backbone clear of traffic in Case D. Filters can be placed on the router 4 port servicing the DECnet host to only allow subnet G traffic. The user side filter would only allow traffic from subnet J. If filters are placed on the server router ports (i.e., subnets H, I, J, or K) it is important to a) maintain them rigorously in the face of moves, adds and changes and b) to remain aware of possible router performance degradation when multiple filters are used.

Summary Most competitive router vendors provide a method for setting filters and taking actions based on matches between filter settings and packets under examination. Research the capabilities of your router(s) and discuss the implementation of filters with your particular vendor representative. They will usually help you out with implementation details, especially if you're in the process of buying one or more routers. Some key features to look for include:

- Types of filters predefined by the vendor
- Filterable items for each protocol
- The ability to construct custom filters
- Number of filters allowed per router port and/or per enabled protocol
- Router performance with filters enabled in a realistic traffic scenario

Important points to remember are that the less filters you use a) the easier the network will be to configure and maintain and b) the lower the possibility of router performance degradation.

- References**
- [1] *Interconnections: Bridges and Routers*, by Radia Perlman, Addison-Wesley, ISBN 0-201-56332-0, 1992.
 - [2] Mogul, J., "Internet Subnets," RFC 917, October 1984.
 - [3] Mogul, J. & Postel, J., "Internet Standard Subnetting Procedure," RFC 950, August 1985.

continued on next page

Establishing Workgroups (*continued*)

- [4] Braden, R. & Postel, J., "Requirements for Internet Gateways," RFC 1009, June 1987.
- [5] *NETWARE, The Professional Reference*, by Karanjit Siyan, New Riders Publishers, 1992.
- [6] *Inside AppleTalk*, Second Edition, by Gurshan Sidhu, Richard Andrews, and Alan Oppenheimer, Addison Wesley, 1993
- [7] "DECnet Routing Layer Functional Specification," April 1988.
- [8] Malamud, C., "DECnet/OSI Phase V: Real OSI or Only Selected Interfaces?," *ConneXions*, Volume 4, No. 10, October 1990.
- [9] Mills, Dave, "Exterior Gateway Protocol Formal Specification," RFC 904, April 1984.
- [10] Hedrick, C., "Routing Information Protocol," RFC 1058, June 1988.
- [11] Moy, J., "OSPF Version 2," RFC 1247, July 1991.
- [12] Baker, F., Coltun, R., "OSPF Version 2 Management Information Base," RFC 1253, August 1991.
- [13] Malkin, G., "RIP Version 2," RFC 1388, January 1993.
- [14] Mogul, J., "Broadcasting Internet Datagrams in the Presence of Subnets," RFC 922.
- [15] Postel, J., "Multi-LAN Address Resolution," RFC 925, October 1984.
- [16] Clark, D., "A Subnetwork Addressing Scheme," RFC 932, January 1985.
- [17] Karels, M., "Another Internet Subnet Addressing Scheme," RFC 936, February 1985.
- [18] Gateway Algorithms and Data Structures Task Force (GADS), "Toward an Internet Standard Scheme for Subnetting," RFC 940, April 1985.
- [19] Hares, S., "Components of OSI: Inter-Domain Routing Protocol (IDRP)," *ConneXions*, Volume 6, No. 5, May 1992.
- [20] *ConneXions*, Volume 3, No. 1, January 1989, "Special Issue: Subnets."
- [21] *ConneXions*, Volume 3, No. 8, August 1989, "Special Issue: Internet Routing."
- [22] *ConneXions*, Volume 5, No. 1, January 1991, "Special Issue: Inter-domain Routing."
- [23] Wolter, M., "Fiber Distributed Data Interface (FDDI)—A Tutorial," *ConneXions*, Volume 4, No. 10, October 1990.

BOB MEINDL holds a B.S. degree from Lafayette College and a M.S. degree from the University of Rhode Island. Since 1992 he has worked for Booz-Allen and Hamilton on a government network design and installation project for a 3,000 node, all fiber network. He is a Member of the IEEE. He can be reached via e-mail as: meindlr@v3.hanscom.af.mil.

Associations: The ATM Forum

by M. Irfan Ali, Newbridge Networks Inc.

Membership

The ATM Forum is an industry wide organization, aimed at promoting ATM within the industry and the end user community. Formed in October 1991, with four members, the ATM Forum currently has approximately 320 members. Of this, over 100 are Principal members and the rest Auditing. The distinction between the two types of membership is reflected in the annual dues charged and the associated privileges.

The success of the ATM Forum is a strong endorsement of ATM as a technology for the future. Bringing together the best points of TDM and Packet Switching, ATM provides a versatile, multifunctional platform that can support a variety of services and traffic types. Besides this, ATM also offers scalability that is difficult to match with other, traditional solutions. It is these capabilities that make ATM popular across a wide spectrum of vendors, carriers, and end users.

Committees

Towards its objective of promoting the ATM technology, the ATM Forum currently has four working committees. The *Technical Committee* is responsible for developing implementation specifications, based primarily on existing standards. To date, the committee has published a *User-Network Interface* (UNI) document that is available to all interested parties at a nominal fee. A more recent version of this document should be available for release in the near future with added functionality on signaling, traffic management, and other related issues. Besides this document, the Technical Committee is also in the process of completing other specifications that should be available shortly.

Besides the Technical Committee, the Forum also has two committees for *Market Awareness and Education* (MA&E); one for North America, and one for Europe. The MA&E committees are responsible for promoting ATM, and the ATM Forum, within the industry. Among several projects undertaken by the MA&E committees on an ongoing basis, a key one is primary market analysis to determine end user requirements and other related information. The MA&E committees are also responsible for making public presentations on ATM and the ATM Forum.

End-user focus

Recognizing the need to keep the end user community closely coupled with the ongoing operation of the ATM Forum, the Forum has recently formed an end user committee. Called the *Enterprise Network Roundtable*, this committee has the charter to bring together interested end users in ongoing discussion on ATM and its implementation. Key among several issues to be discussed are, interoperability between different vendors and a graceful migration of the installed base to ATM.

M. IRFAN ALI is Assistant Vice-President, ATM Product Marketing with Newbridge Networks, Inc. Prior to joining Newbridge, Irfan was with Northern Telecom. He started his career with Northern Telecom as a part of the Strategic Design and Product Evolution group at Bell-Northern Research (BNR: the R&D subsidiary of Northern Telecom). As a part of this group, he was involved in several areas of high speed networking, including Frame Relay, Wideband Switching, ATM, and SONET. After 4 years at BNR, he joined Northern Telecom as a part of the Carrier Systems division. His primary responsibility was market development for strategic technology. The key areas of focus were broadband networking and wireless communications. Irfan was also a member of the founding team for the ATM Forum and is currently on the Board of Directors for the Forum as Vice-President Marketing. He holds a masters in Electrical Engineering from Southern Methodist University, Dallas, Texas and is currently completing his MBA from the same school.

Associations: The Frame Relay Forum

by Alan Taffel, Alcatel Data Networks

Group charter

The Frame Relay Forum is an association of vendors, carriers, users, and consultants committed to the implementation of Frame Relay in accordance with national and international standards. The group was formed in 1991 and was the first of its kind in the world. It now maintains chapters in North America, Europe, Australia/NZ, and Japan.

The Forum's technical committees take existing standards, which are necessary but not sufficient for full interoperability, and create *Implementation Agreements* (IAs). These IAs represent an agreement by all members of the Frame Relay community as to the specific manner in which standards will be applied, thus ensuring interoperability. At the same time, the Forum's marketing committees are chartered with worldwide market development through education as to the benefits of Frame Relay technology.

Membership

The Frame Relay Forum supports two primary types of members: *Worldwide Members* and *Auditing Members*. Worldwide Members, sometimes known as Full Members, enjoy all of the benefits of the Forum: access to all Forum marketing and technical materials, listing as a Member in Forum documents, free attendance for two at the Annual and General Meetings, unlimited participation in Committee meetings, and voting privileges on technical meetings, and Forum elections. Auditing Members receive the bulk of these benefits, but do not have voting rights and only one free attendance to the Annual and General Meetings. Additional members can also attend these meetings by paying a nominal admission fee. The Forum has over 115 Worldwide Members as well as dozens of Auditing Members.

Committees

Each Forum chapter maintains two committees in support of Forum goals. The *Technical Committees* in each chapter work in close cooperation to take existing national and international standards and works to create common Implementation Agreements (IA) which define how those standards will be applied. The Committee typically maintains several Sub-Committees which work on specific IAs. Once a draft IA has been produced and agreed to by the Technical Committee, it is circulated to all Worldwide Members for ratification.

The *Marketing Committees* create materials and organize activities which promote the implementation of Frame Relay in their respective territories. Materials include newsletters, application guides, and slide presentations. Activities may include presentation seminars, trade show participation, speakers bureau programs, or user roundtables. Each chapter determines the materials and activities which best meet the needs of their region.

ALAN TAFFEL is the president and chairman of the Board of the Frame Relay Forum. He is also Vice President, Marketing for Alcatel Data Networks, which is based in Reston, Va. Prior to his tenure at Alcatel, Mr. Taffel held several executive positions within Sprint Corporation, a charter member of the Frame Relay Forum. He holds a Bachelor of Sciences degree in Information and Computer Sciences from the University of California, where he graduated with honors.

The Secretariats for all 3 associations are located in Mountain View, CA. Those interested in membership, meeting schedules, or in ordering marketing or technical materials should contact:

The ATM Forum:	+1 415-962-2585
The Frame Relay Forum:	+1 415-962-2579
The SMDS Interest Group:	+1 415-962-2590

Associations: The SMDS Interest Group

By Robyn Aber, Bellcore

Group charter

The *SMDS Interest Group* (SIG) is an industry association of local and long distance telecommunications carriers, network and computing equipment manufacturers, consultants, and end-users committed to the advancement of worldwide *Switched Multimegabit Data Service* (SMDS) as an open, interoperable solution for high performance data connectivity. By advancing public understanding of SMDS, and by developing specifications that support compatibility among different networks and equipment, the SIG expects to further the development of SMDS as a vital part of our national telecommunications infrastructure, to speed product development and deployment, to foster the exchange of technical information among industry members, and to ensure the worldwide interoperability of network systems.

Membership

Companies that participate in the SIG range from small consultancies to corporations that appear on every list of the nation's largest and most influential firms, and members include many international enterprises. From its five founding members in February 1991, the SIG has grown to more than sixty corporate members and around fifty users who are involved in the *SMDS Users Group*. SIG members also work closely with their overseas counterparts, the *European SMDS Interest Group* (ESIG) and the *Pacific Rim SMDS Interest Group*.

Working Groups and their activities

Technical Working Group (TWG) members initiate work on specs that support compatibility among different networks and network equipment. They have developed specifications for SMDS to support common protocol architectures like TCP/IP, Novell, AppleTalk, DECnet, SNA, and OSI. They have also specified the means for SMDS to transparently interoperate with frame-based and cell-based technology platforms. In early 1993 this group approved specs for frame-based access to SMDS and Frame Relay access to SMDS; both of which enable access at 56 Kbps, 64 Kbps, and increments of Nx56/64 Kbps, so that users running applications at lower speeds can make use of SMDS' connectivity features. The TWG is currently finalizing an interface spec to allow frame relay frames to be encapsulated within an SMDS interface protocol for transport over SMDS networks.

Intercarrier Working Group (IWG) members are carriers who support interoperability between local and long distance service offerings, and who define network management as it relates to the interoperability of SMDS. This group recently approved two documents that cite essential elements for the deployment and interconnection of local area and interexchange SMDS networks, and that describe "tools" for the operations management of intercarrier SMDS. The net impact of these guiding principles is that they will extend the reach of SMDS to national and global service interconnection capability.

Public Relations & User Awareness Working Group members advance public awareness and understanding of SMDS and its applications. This group organizes seminars, sponsors conference sessions, organizes trade show demonstrations, manages a speakers' bureau, writes articles and newsletters, and offers tutorials in a variety of formats that explain SMDS capabilities and applications.

ROBYN ABER serves as treasurer of the SMDS Interest Group and chairs its working group for public relations and user awareness. She is also a fast packet product manager at Bellcore.

SMDS Users Group members are current and potential customers of SMDS that provide a clearinghouse for the dissemination of SMDS technical reports and user experiences. The group promotes open discussion of end-user experiences and issues with SMDS, and participates in the technical evaluation and evolution of SMDS and related technologies.

Letters to the Editor

In response to Craig Partridge's review in the May, 1993 issue:

Thank you for your favorable review of my book, *TCP/IP: Architectures, Protocols, and Implementation* (McGraw-Hill).

Regarding errors

There were a few misconceptions among your statements about Errors. First of all, the sample code was not meant to be just read, but *typed in, compiled, run, and then expanded* as students wish. The server does not bind to a well-known port because end-user code had better take whatever port it can get! Note that the server prints the port that it gets, so that the user can feed this to each client as a parameter. (The client software carries this through.) Obviously, it would be very easy to alter this code to run at a well-known port, if it were used as the basis of a new system service.

Karn/Partridge Algorithm

With regard to Karn's algorithm (which should be the Karn/Partridge algorithm), the discussion in the book was limited to what TCP *does*, as opposed to *why*. This was a conscious simplification, meant to make it easy for users new to TCP/IP to understand what is going on, but probably very irritating to experts—and especially to one of the authors of the algorithm! In the Karn/Partridge paper: "Improving Round-Trip Time estimates in Reliable Transport Protocols," you state:

1. "When an acknowledgement arrives for a packet that has been sent more than once ... ignore any round-trip measurement based on this packet ..."
2. "*In addition, the backed off RTO*" (retransmission time out) "*for this packet is kept for the next packet. Only when it ... is acknowledged without an intervening retransmission will the RTO be recalculated from SRTT.*"

The italics are the authors'. However, I should have mentioned that the interpretation of a late or missing acknowledgement as congestion was due to Van Jacobson, and a fuller discussion the logic in the Karn/Partridge paper will be included in future editions.

E-mail

With regard to mail exchangers and e-mail gatewaying, page 291 of the book does point out: "A mail exchanger also can act as a gateway to external non-Internet style mail services...." I would like to include a lot more detail about DNS servers and Resource Records in a future edition. The effective operation of the distributed DNS database is one of the Internet miracles.

Layering

My statement on page 26 regarding layering was: "One motivation for following a layered model is to give communications software a structure that is rational, simple, and easy to modify,"—a mild statement. I believe that positive points should have been awarded for (contrary to common custom): 1) not mis-describing the OSI Presentation layer 2) not describing the OSI layers at all! However, the requisite incantation of the OSI model gets on my nerves too, and perhaps a future edition will include the truth about layering, for example:

1. The most important layer interface is the production of common Application Programming Interfaces. OSI refused to look at this for over a decade.
2. Don't move data between internal layers. Don't take headers apart. So-called service interfaces are fictions that indicate the *minimum* of information that needs to be shared.

3. There are less than ten people in the world who understand the OSI presentation layer.
4. The OSI session, presentation, and application layers are not layers.

References

- [1] P. Karn & C. Partridge, "Improving Round-Trip Time estimates in Reliable Transport Protocols," Proceedings of ACM SIGCOMM-87, 1987.
- [2] V. Jacobson, "Congestion Avoidance and Control," Proceedings of ACM SIGCOMM-88, 1988.

—Dr. Sidnie Feit
The Standish Group

Dear Ole,

OSI debates

I continue to enjoy reading the OSI debates. I categorize the contributions as follows:

(1) *Open minded contributions*: an example is the "The IETF integrates OSI related work" by Erik Huizer, SURFnet [*ConneXions*, Volume 7, No. 6, June 1993, p 26–28]. Open minds try to use the best of both worlds to the user's benefit. The problem with open minds is that they have to try hard to carry conviction. And they have to work hard to understand the pros and cons of more than one way to Rome.

(2) *Close minded contributions*: an example is "Network Management: Status and Challenges" by Marshall T. Rose, Dover Beach Consulting, Inc. [*ConneXions*, Volume 7, No. 6, June 1993, p 11–17]. Closed minds always succeed in reducing complex matter to a size which is clear and manageable. They are the heroes. But they always run the danger of missing something which happens outside of their view of the world.

Public data networks

(3) *Contributions which show an opening mind*: an example is "You cannot promote OSI Applications over OSI networks" by Paul Barker, UCL and Colin Robbins, NeXor Ltd. [*ConneXions*, Volume 7, No. 5, May 1993, p 16–22]. Opening minds show their capability to learn. We all tend to neglect that besides OSI and TCP/IP there is CCITT and the world of public data networks. All of these protocols and networks come down to interface connectors—this is where the challenge really is.

Which category is best? Do the OSI debates indicate a technical problem or a political one? Do we know where we want to go?

—Dr. Harald Hoffmann
OSIconsult Kommunikationssysteme GmbH, Wien, Austria

House Bill HR1757 Available on line

The revised version of HR1757 (see *ConneXions*, Volume 7, No. 8, August 1993), the *National Information Infrastructure Bill* introduced by Rep. Boucher is now available on the CPSR Internet Library. FTP/WAIS/Gopher to Internet host cpsr.org and retrieve the file /cpsr/nii/hr1757_july_1993.txt. If you do not have FTP/WAIS/Gopher, e-mail listserv@cpsr.org with the word "help" as the body of the text.

—Dave Banisar
CPSR Washington Office

Book Review

UNIX, POSIX, and Open Systems: The Open Standards Puzzle by John S. Quarterman and Susanne Wilhelm, Addison-Wesley, 1993 (ISBN 0-201-52772-3)

Background

Users like open systems because they promise freedom from a single vendor's control, competition among vendors, and, most importantly, lower prices. Vendors like open systems because—well, most vendors don't really like open systems much at all, for pretty much the same reasons that users like them. A good deal of the open systems movement has focused around the UNIX operating system, since it has been the only powerful, widely used system available from a number of different vendors. POSIX was created in an attempt to standardize various aspects of UNIX among these vendors. This detailed and quite complete book, written by two long-time participants in the POSIX work, describes these standards and many more.

Organization

Beginning with a general discussion of computing, open systems, and open standards, the book moves on to a description of formal standards bodies, industry organizations, and user groups, describing the role and process for each. There is then a large section on POSIX and its standards, including discussions of issues like internationalization and how POSIX was affected by the perceived requirement for language independence. The book closes with a look at some of the many standards-based profiles that have been defined for various areas in computing. A whimsical appendix offers a word search puzzle, challenging the reader to find the hidden acronyms for many of the organizations and specifications discussed earlier.

Descriptions of international standards bodies like ISO are not hard to find. While this book (quite appropriately) reiterates some of this widely available material, its greatest contribution is a deep explanation of the POSIX standards and the POSIX process. Neither of these is simple, and their discussion takes up a large share of the book. This emphasis is also apparently due to the authors' greater familiarity with this area—they have chosen to write primarily about topics they truly understand, an approach one hopes more writers will adopt.

UNIX

UNIX's status as the only powerful system available on many vendors' hardware is at an end, however. Microsoft's Windows NT is here, and it promises to give UNIX a run for its money, in the most literal sense. While those selling UNIX have been working through the bureaucratic morass described so well in this book, Microsoft has retained a single point of control over NT. Rather than relying on the slowly forged consensus of POSIX committees, Microsoft can enhance their system as they choose, based entirely on what really matters: input from their users. While this monopolistic approach doesn't sit well with vendors—they want something at least a little unique to sell—it is exactly what users want.

One could argue that UNIX users (although probably not most UNIX vendors) would be better served if System V's new owner Novell begins to behave more like Microsoft and less like they're in the open systems business. After all, Microsoft's MS-DOS, for all its faults, has given users exactly what they want: competition and very low prices for hardware, and a rich suite of competing applications. Isn't this exactly what open systems were supposed to do? True, the Microsoft approach has had the side effect of making Bill Gates the richest man in America, but only his competitors care about that; users certainly don't.

Excellent reference

That said, it is far too soon to write an obituary for UNIX or for the open systems movement as it stands today. For the large number of people who work in this world, this book provides an excellent reference to an important class of standards. The authors have made every effort to create an interesting, readable text, and have in the main been quite successful (although the very nature of the book's subjects can occasionally leave one feeling as if one were reading some high-tech version of *Roberts Rules of Order*). For reference or answers to specific questions, however, there's nothing like it available.

—David Chappell

Networking Papers Available

The following paper is available via anonymous FTP to host `ftp.ee.lbl.gov`; retrieve `WAN-TCP-models.{1,2}.ps.Z`. Each file uncompresses to a bit over 500KB of *PostScript*. The title of the paper is "Empirically-Derived Analytic Models of Wide-Area TCP Connections: Extended Report."

Abstract

We analyze 2.5 million TCP connections that occurred during 14 wide-area traffic traces. The traces were gathered at five "stub" networks and two internetwork gateways, providing a diverse look at wide-area traffic. We derive analytic models describing the random variables associated with Telnet, NNTP SMTP, and FTP connections, and present a methodology for comparing the effectiveness of the analytic models with empirical models such as *tcplib*. Overall we find that the analytic models provide good descriptions, generally modeling the various distributions as well as empirical models and in some cases better.

A (considerably shorter) companion paper is also available from the same host as `WAN-TCP-growth-trends.ps.Z`. The title of the second paper is "Growth Trends in Wide-Area TCP Connections."

Abstract

We analyze the growth of a medium-sized research laboratory's wide-area TCP connections over a period of more than two years. Our data consisted of six month-long traces of all TCP connections made between the site and the rest of the world. We find that SMTP, FTP, and X11 traffic all exhibited exponential growth in the number of connections and bytes transferred, at rates significantly greater than that at which the site's overall computing resources grew; that individual users increasingly affected the site's traffic profile by making wide-area connections from background scripts; that the proportion of local computers participating in wide-area traffic outpaces the site's overall growth; that use of the network by individual computers appears to be constant for some protocols (Telnet) and growing exponentially for others (FTP, SMTP); and that wide-area traffic geography is diverse and dynamic.

—Vern Paxson
Systems Engineering
Lawrence Berkeley Laboratory
vern@ee.lbl.gov
(510) 486-7504

Call for Papers

The *First International Symposium on Interworking* was held in Bern during 1992 and attracted a large number of contributions and participation of delegates from around the world. The Second International Symposium on Interworking, *INTERWORKING '94*, is scheduled for May 4–6, 1994 in Sophia Antipolis, France. Hosted by France Telecom and AscomTech in Switzerland, the symposium will highlight the importance of interworking and interoperability.

Rationale for the symposium

Telecom operators have announced their plans for ATM field trials and the ATM Forum aims for the user awareness and application of ATM in private networks. It is a common view that ATM will offer new opportunities for global multimedia communications. Initial applications will be driven from the interconnection of all types of LANs and MANs. A full integration will lead to what is known as the *Integrated Broadband Communication Network* (IBCN). This will require an extended provision of interworking between the heterogeneous infrastructures involved. The interworking facilities thus play a very important role in the provision of world-wide communications.

Topics

This Symposium provides the platform for exchange of views on heterogeneous network evolution, concepts, services, equipment and user requirements. It will highlight the importance of interoperability between equipment, protocols, signalling and network management for the provision of end-to-end services at an international level. A keynote speech will address the future focus on research and development of advanced networks under the fourth framework programme in Europe. Areas to be addressed by contributed papers include:

- Network architectures, interworking principles and evolution
- ITU-TSS, ISO, ETSI, ANSI, IEEE standards for interworking
- Heterogeneous Network Architectures and Scenarios
- Protocol Converters, Bridges, Routers, Brouters, Hubs
- LAN–MAN–WAN interworking
- Service, Network, and Media Adaptation Systems
- Interoperability in public and private networks
- Signalling, Management and Control aspects
- Multimedia Service Support
- Models, Simulation and Performance Analysis
- Test beds, Field Trials and Pilot Projects experiences
- Market trends, new products and systems

Submissions

A 500 word abstract should be sent to:

Dr. S. Rao, Technical Committee chairman
AscomTech
Freiburgstraße 370
CH-3018 Berne
Switzerland
Tel: +41 31 999 4397
Fax: +41 31 991 5211
E-mail: rao@tech.ascom.ch

Important dates

Abstracts due:	September 15, 1993 (send it <i>today</i> !)
Notification of acceptance:	November 15, 1993
Full papers due:	January 15, 1994

Call for Papers

The *5th IFIP Conference on High Performance Networking* will be held in Grenoble, France, June 27–July 1, 1994.

Topics

This workshop belongs to the series started in 1987 in Aachen, followed by Liege in 1988, Berlin in 1991 and Liege in 1992. It aims at presenting and discussing evolution in the framework of high-speed networking and computing in private and public networks. Original contributions on the following topics are solicited:

- *New MAC Services and Protocols*: Gigabit networks, ATM-based systems
- *Enhanced Network and Transport Services and Protocols*: Multi-peer services and protocols, Admission and congestion control, Time-constraint management
- *New Services and Protocols*: Synchronization semantic and management, Protocols for groupware communication, Video over high speed networks, Quality of Service semantics
- *New Applications*: Multimedia, Distribution network algorithms, Groupware communication
- *Internetworking*: Routing in high performance multimedia networks, Bridges and routers technology and protocols, Meshed architectures
- *Implementation and Performance Evaluation*: MAC Performance in high speed networks, Efficient Protocol Implementation

Venue

The conference will be organized by the IMAG Institute and IBP Institute, and will be held in Grenoble. Grenoble is located one hundred kilometers from Lyon and one hundred and fifty kilometers from Geneva. Grenoble is the heart of the French Alps. Following the 1968 Olympic Games, Grenoble developed a high technology R&D park around one of the most famous French Universities.

Submissions

Papers must be written in English and should not exceed 12 pages single-spaced, or 20 pages double-spaced. The front page should contain author names, addresses, phones, faxes, and e-mails, as well as a 150 words abstract. Papers that scope with the topics will be refereed. Suggestions for half- or full-day tutorials are also welcome. (Tutorials will be organized on June 27 and 28, 1994.) Authors of accepted papers will be requested to sign a copyright release from IFIP. A participant edition of the proceedings will be made available at the conference from the camera-ready copy which will be used later for the publication of the proceedings by Elsevier/North Holland. Four copies of the paper should be submitted to:

Serge Fdida
 Université René Descartes–UFR Maths-Info
 Laboratoire MASI
 45, rue des Saints-Pères
 75006 Paris, France
 Phone: +33 (1) 42 86 21 36 • Fax: +33 (1) 42 86 22 31
 E-mail: fdida@masi.ibp.fr

Important dates

Today:	Notification of Intent to Submit a Paper
October 30, 1993:	Full Paper Submission Deadline
January 31, 1994:	Notification of Acceptance
March 31, 1994:	Camera-Ready Copy Due

Call for Papers

"Signaling Protocols and Services for Broadband ATM Networks" will be the focus of a Special issue of the *International Journal of Digital and Analog Communication Systems* (IJDACS) published by John Wiley and Sons, Ltd. (The issue will be Volume 7, No. 2, June 1994.)

Background

Broadband networks based on ATM provide the telecommunications world with a very flexible transport and switching infrastructure. However, the success of an ATM-based network will also be dependent upon the availability of a wide variety of attractive services, most of which have not yet been defined. Therefore, a flexible and extensible control environment is required to take full advantage of the potential provided by the transport network.

In a first phase, signaling protocols will be based on extensions to existing narrow band control protocols Q.931 (at the UNI) and perhaps ISUP (at the NNI). However, new multimedia, multiparty services, possibly requiring distributed processing, will pose a formidable challenge for the definition of new control protocols. In addition, the fast development of extensive private LAN and MAN networks, interconnected to the public networks, is accelerating the integration of data communications services (mostly connectionless) with voice and video services (mostly connection oriented). Multiparty applications, such as videoconferencing and shared work spaces, will require negotiation and transaction processing systems to control special resources such as video bridges and format converters. Trade-offs between "network oriented" and "terminal oriented" approaches will have to be examined. Moreover, evolution strategies will have to be elaborated in order to allow interworking with earlier protocols and across heterogeneous networks. The concept of separating "Call (or Session) Control" from "Connection (or Resource) Control" will provide mechanisms for avoiding unnecessary preallocation of resources, as the Call Control protocols negotiate service aspects edge-to-edge in a technology-independent framework prior to convoking any physical resources. Future signaling protocols will have to also provide support for mobile users, personalized user profiles and diverse terminal capabilities. Finally, interfaces have to be foreseen to allow interworking with network management and operations support systems, such as described in IN and TMN models.

Topics

As a guide to the authors considering a contribution to this special issue, the following keywords are provided: Call Control; Connection Control; Control Network Architectures; Distributed Processing Systems; Information Networking Systems; Intelligent Networks; Management of Network Services; Multimedia Applications; New Service Applications; Object-oriented Protocols; Personalized Communications; Presentation Control; Protocol Architectures; Protocol Evolution & Interworking; Signaling Requirements; Supplementary Services; Transaction Processing Systems.

Submissions

Papers that address the above issues, including modeling, specification, simulation, prototyping, experimentation and standards activities, will be considered and should be submitted by October 15, 1993 to one of the following editors:

Dr. Martin De Prycker
Alcatel Bell Telephone
Bell Telephone Man. Co.
Francis Wellesplein 1
B-2018 Antwerpen
BELGIUM

Dr. Mario P. Vecchi
Bellcore, MRE-2Q268
445 South Street
Morristown, NJ 07962
USA
mpv@faline.bellcore.com

Call for Papers

INET '94, the Annual Conference of the Internet Society will be held in conjunction with the 5th Joint European Networking Conference (JENC5), in Prague, Czech Republic, June 13–17, 1994.

Topics

The Internet Society (ISOC) and Réseaux Associés pour la Recherche Européenne (RARE) have joined forces to organize a global networking conference. Focusing on worldwide issues of computer-based networking, the goal of the conference is to bring together individuals from academia, industry and government who are involved with planning, developing, implementing, managing, funding, and using national, regional and international research, academic, and commercial computer networks. The official language of the conference is English. Possible topics for paper submission include but are not limited to the following:

- *Network Technology:*
 - Progress towards international open network protocols
 - Security, management and authentication in managing networks
 - Transmission, routing, and transport technologies
 - Technologies of the '90s and 21st century
 - Very high speed networks
 - LAN/WAN integration and interworking
 - Support for mobility
- *Network Engineering and Operation:*
 - Application of network technology to provide networking services
 - Interoperability among national and international networks
 - Network management systems and methods
 - Reliability and performance engineering
 - Issues related to scaling
 - Emergency response organizations and support
 - Resource allocation and control
- *Distributed Applications and Their Enabling Technologies:*
 - Collaboration technologies
 - Multimedia issues
 - Mail and directory services
 - Workstation teleconferencing
 - Computer supported collaborative work
 - Interoperability of application services
 - Management protocols, systems and methods
 - Security aspects of distributed applications
 - Distributed application development environments
 - Visions for the future of internationalized services
- *Support and Training for International Communities of Interest:*
 - Support of international collaboration
 - Globalization of user support services

continued on next page

Call for Papers (*continued*)

Access to scientific papers and data across national boundaries
 Supercomputing
 High energy physics, atmospheric modeling
 Scientific applications
 Education/distance learning
 Medical research and clinical applications
 Libraries

- *Work and play in Cyberspace:*

Networking and the arts
 High payoff application areas
 Tools for user education and training

- *User Information:*

Networked information retrieval
 User documentation
 Navigation services
 Document delivery services
 Library services

- *Regional Issues:*

Africa
 Asia-Pacific Rim
 Former Soviet Republics
 Latin America
 North America
 Western and Central Europe

- *Policy Issues:*

Globalization of services
 Commercialization, privatization and public access
 Coordination of international resources
 Copyright and intellectual property rights
 Appropriate use and speech restrictions
 International security policy
 Privacy and data protection
 Telecommunications policy
 Coordination: users/providers and suppliers/policy makers
 Interworking issues: commercial and academic network providers

Submissions

All papers must be written in English. Electronic submission is highly recommended. ASCII or uuencoded *PostScript* can be sent by e-mail to `inet-jenc-submit@rare.nl`

PostScript documents can be sent via anonymous FTP to Internet host `erasmus.rare.nl` (IP address 192.87.30.2), into the directory `pub/inet-jenc/submit`. Please note that files deposited in this directory can only be written once and cannot be deleted afterwards.

Should electronic submission be impossible, please submit 6 copies of double-spaced full paper manuscript (maximum of 4000 words) with an abstract to the INET-JENC Secretariat at the address given below.

Demonstrations and presentations

There will be the opportunity for participants to present their projects or activities in the form of a demonstration. Proposed demonstrations should be documented with a description not exceeding one page. There will be provisions for presentations describing the activities in RARE Working Groups and IETF Areas.

Important dates

December 15, 1993:	Full manuscript due
December 15, 1993:	Proposals for demonstrations due
March 1, 1994:	Notification of acceptance to authors
April 11, 1994:	Camera-ready papers due

Publication

Conference proceedings containing full papers will be handed out to the participants. A selection of the best papers will be published as a special issue of *Computer Networks and ISDN Systems*.

Workshop and tutorials

A workshop on the installation and use of networking technology is planned to take place adjacent to the conference week. Travel and tuition support may be available for selected attendees. Additional information will be forthcoming. Tutorials are planned to be held on June 13 and 14, 1994.

Points of contact

Conference Chair :	Geoff Manning	GM1@ib.rl.ac.uk
Program Chair :	Bernhard Plattner	plattner@komsys.tik.ethz.ch
Program Chair Deputy:	Hannes P. Lubich	lubich@komsys.tik.ethz.ch
Local Organization:	Jan Gruntorad	tkjg@earn.cvut.cs
ISOC Liaison :	Larry Landweber	landweber@cs.wisc.edu
RARE Liaison:	Tomaz Kalin	kalin@rare.nl

General inquiries

To be added to the conference e-mail distribution list send a message to:

Internet: inet-jenc-request@rare.nl

X.400: C=nl; ADMD=400net; PRMD=surf;
O=rare; S=inet-jenc-request

For further information contact:

INET-JENC Secretariat
c/o RARE Secretariat
Singel 466-468
NL-1017 AW Amsterdam
The Netherlands
Tel.: +31 20 639 1131
Fax: +31 20 639 3289
Internet: inet-jenc-sec@rare.nl
X.400: C=nl; ADMD=400net; PRMD=surf;
O=rare; S=inet-jenc-sec

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute



Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years
International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to: CONNEXIONS

Back issues available upon request \$15./each
Volume discounts available upon request

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779
connexions@interop.com

CONNEXIONS